



(12) **United States Patent**
Resch et al.

(10) **Patent No.:** **US 10,841,081 B2**
(45) **Date of Patent:** ***Nov. 17, 2020**

(54) **THRESHOLD OBLIVIOUS PSEUDORANDOM FUNCTION IN A KEY MANAGEMENT SYSTEM**

(58) **Field of Classification Search**
USPC 380/44
See application file for complete search history.

(71) Applicant: **International Business Machines Corporation**, Armonk, NY (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

(72) Inventors: **Jason K. Resch**, Chicago, IL (US); **Hugo M. Krawczyk**, Tarrytown, NY (US); **Patricia Sagmeister**, Adliswil (CH); **Martin Schmatz**, Rueschlikon (CH); **Mark D. Seaborn**, Algonquin, IL (US)

5,991,414 A 11/1999 Garay et al.
6,118,874 A * 9/2000 Okamoto H04L 9/0894
380/282
8,468,368 B2 6/2013 Gladwin et al.
8,538,029 B2 9/2013 Li et al.
9,049,011 B1 * 6/2015 Agrawal H04L 9/085
(Continued)

(73) Assignee: **INTERNATIONAL BUSINESS MACHINES CORPORATION**, Armonk, NY (US)

OTHER PUBLICATIONS

National Institute of Standards and Technology; The NIST Definition of Cloud Computing; Special Publication 800-145; Dec. 13, 2016; pp. M-7 to M-13.

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 161 days.

This patent is subject to a terminal disclaimer.

Primary Examiner — Amare F Tabor

(74) *Attorney, Agent, or Firm* — James Nock; Andrew D. Wright; Roberts Calderon Safran & Cole, P.C.

(21) Appl. No.: **15/979,937**

(57) **ABSTRACT**

(22) Filed: **May 15, 2018**

A computing device is configured to divide an Oblivious Pseudorandom Function (OPRF) key to generate a plurality of N partial keys, distribute a respective one of the plurality of N partial keys to a corresponding plurality of N Key Management System (KMS) units. The computing device receives from a threshold number T of KMS units, a plurality T partial blinded keys, wherein the plurality T partial blinded keys are based on processing of a value of a blinded key received by a respective KMS unit and a corresponding stored partial key of the N partial keys, combines the plurality T of partial blinded keys into the blinded key, processes the blinded key based on the blinding key in accordance with an OPRF unblinding operation to generate a key and accesses secure information based on the key.

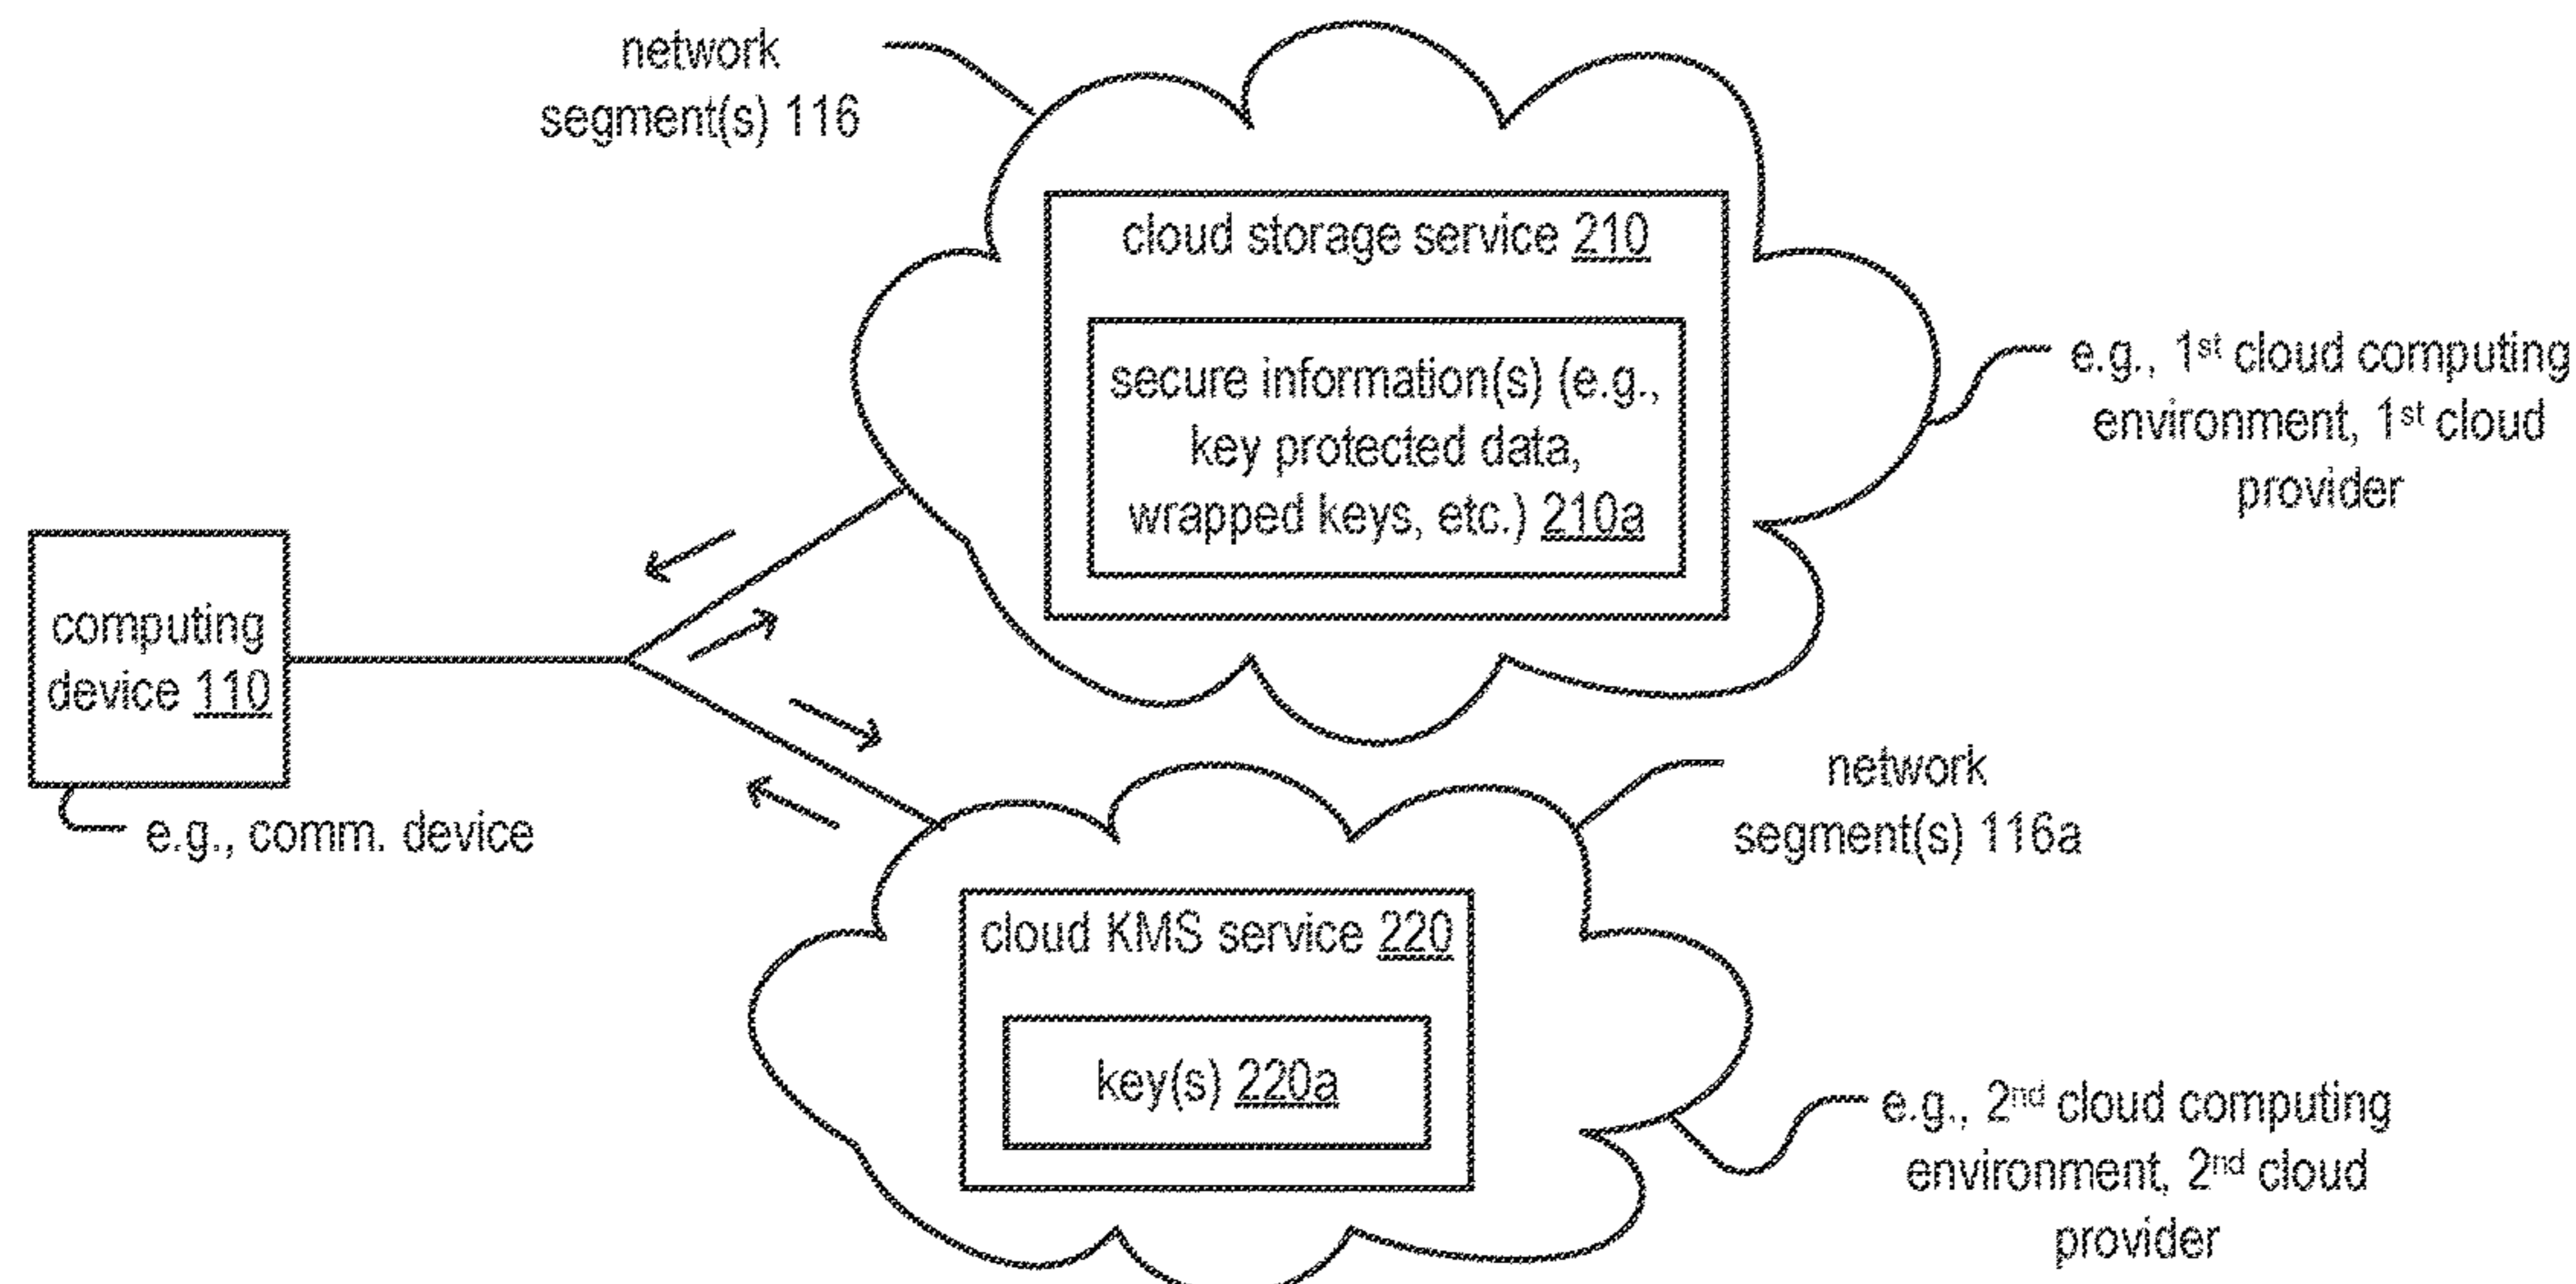
(65) **Prior Publication Data**

US 2019/0356475 A1 Nov. 21, 2019

20 Claims, 10 Drawing Sheets

(51) **Int. Cl.**
H04L 9/08 (2006.01)
H04L 9/06 (2006.01)
H04L 29/06 (2006.01)

(52) **U.S. Cl.**
CPC **H04L 9/0662** (2013.01); **H04L 9/0819** (2013.01); **H04L 63/062** (2013.01); **H04L 2209/04** (2013.01)



(56)

References Cited

U.S. PATENT DOCUMENTS

9,380,036	B2 *	6/2016	Parann-Nissany	H04L 63/06
9,380,037	B2 *	6/2016	Parann-Nissany	H04L 63/06
2011/0055585	A1 *	3/2011	Lee	H04L 9/0844 713/183
2011/0126295	A1	5/2011	Resch	
2012/0121080	A1	5/2012	Kerschbaum	
2017/0177899	A1	6/2017	Ng et al.	
2017/0223008	A1 *	8/2017	Camensch	H04L 9/0894
2018/0076956	A1 *	3/2018	Camensch	H04L 9/0894
2018/0375652	A1 *	12/2018	Karame	H04L 9/0819

* cited by examiner

FIG. 1A




FIG. 1B




FIG. 1C



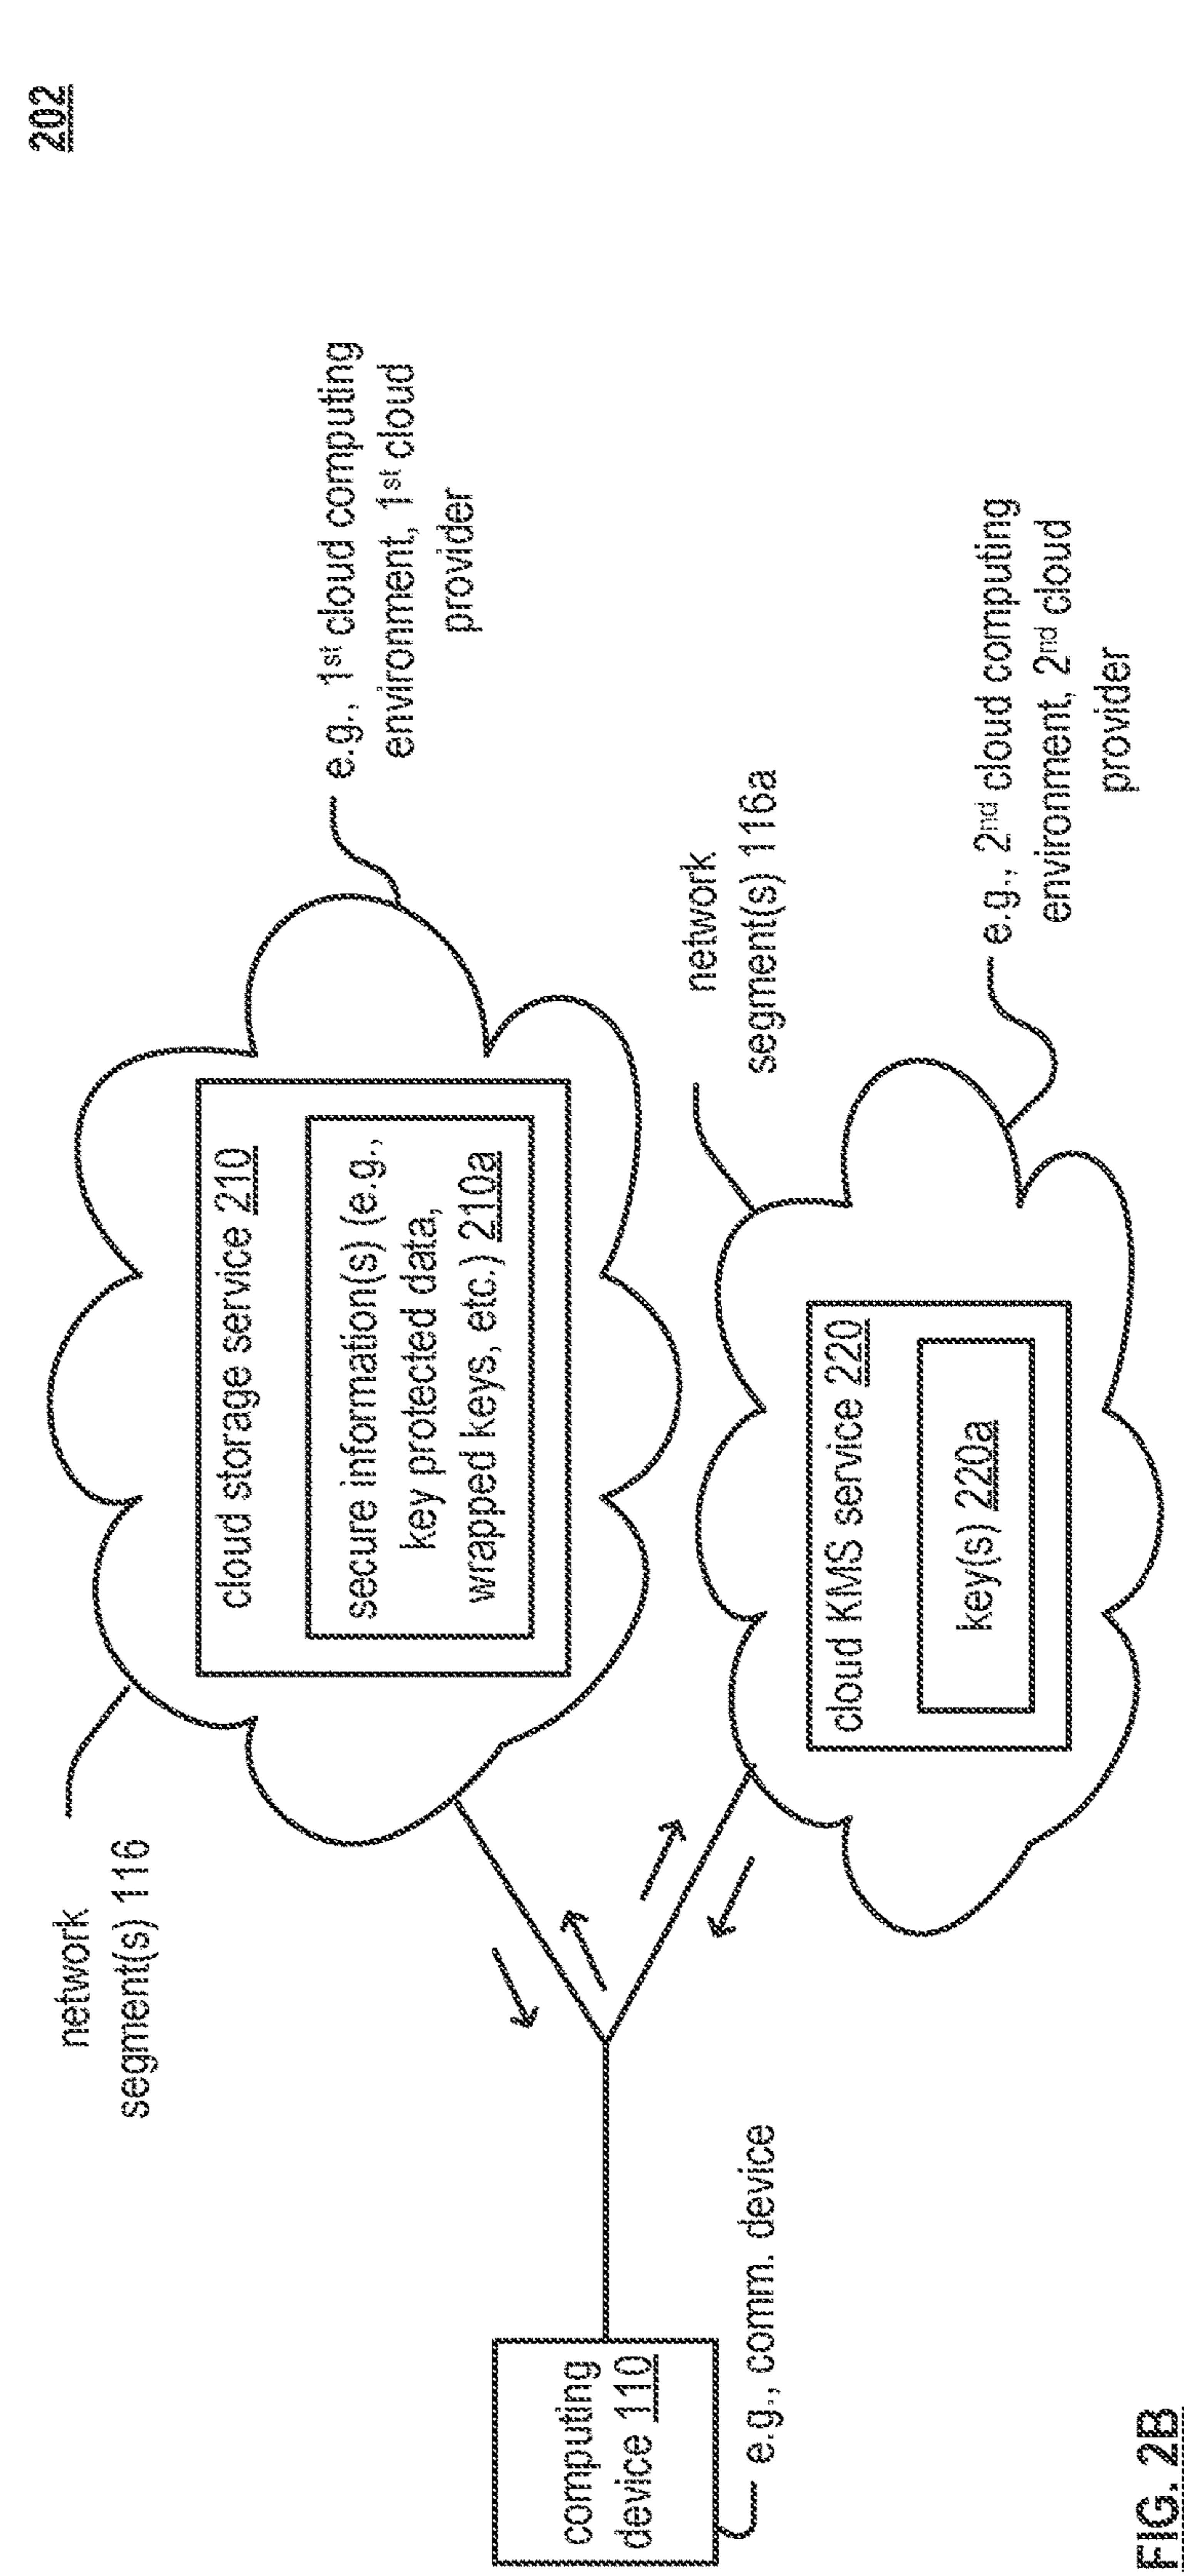



FIG. 1D








FIG. 3A




FIG. 3B




FIG. 4A




FIG. 4B




FIG. 4C

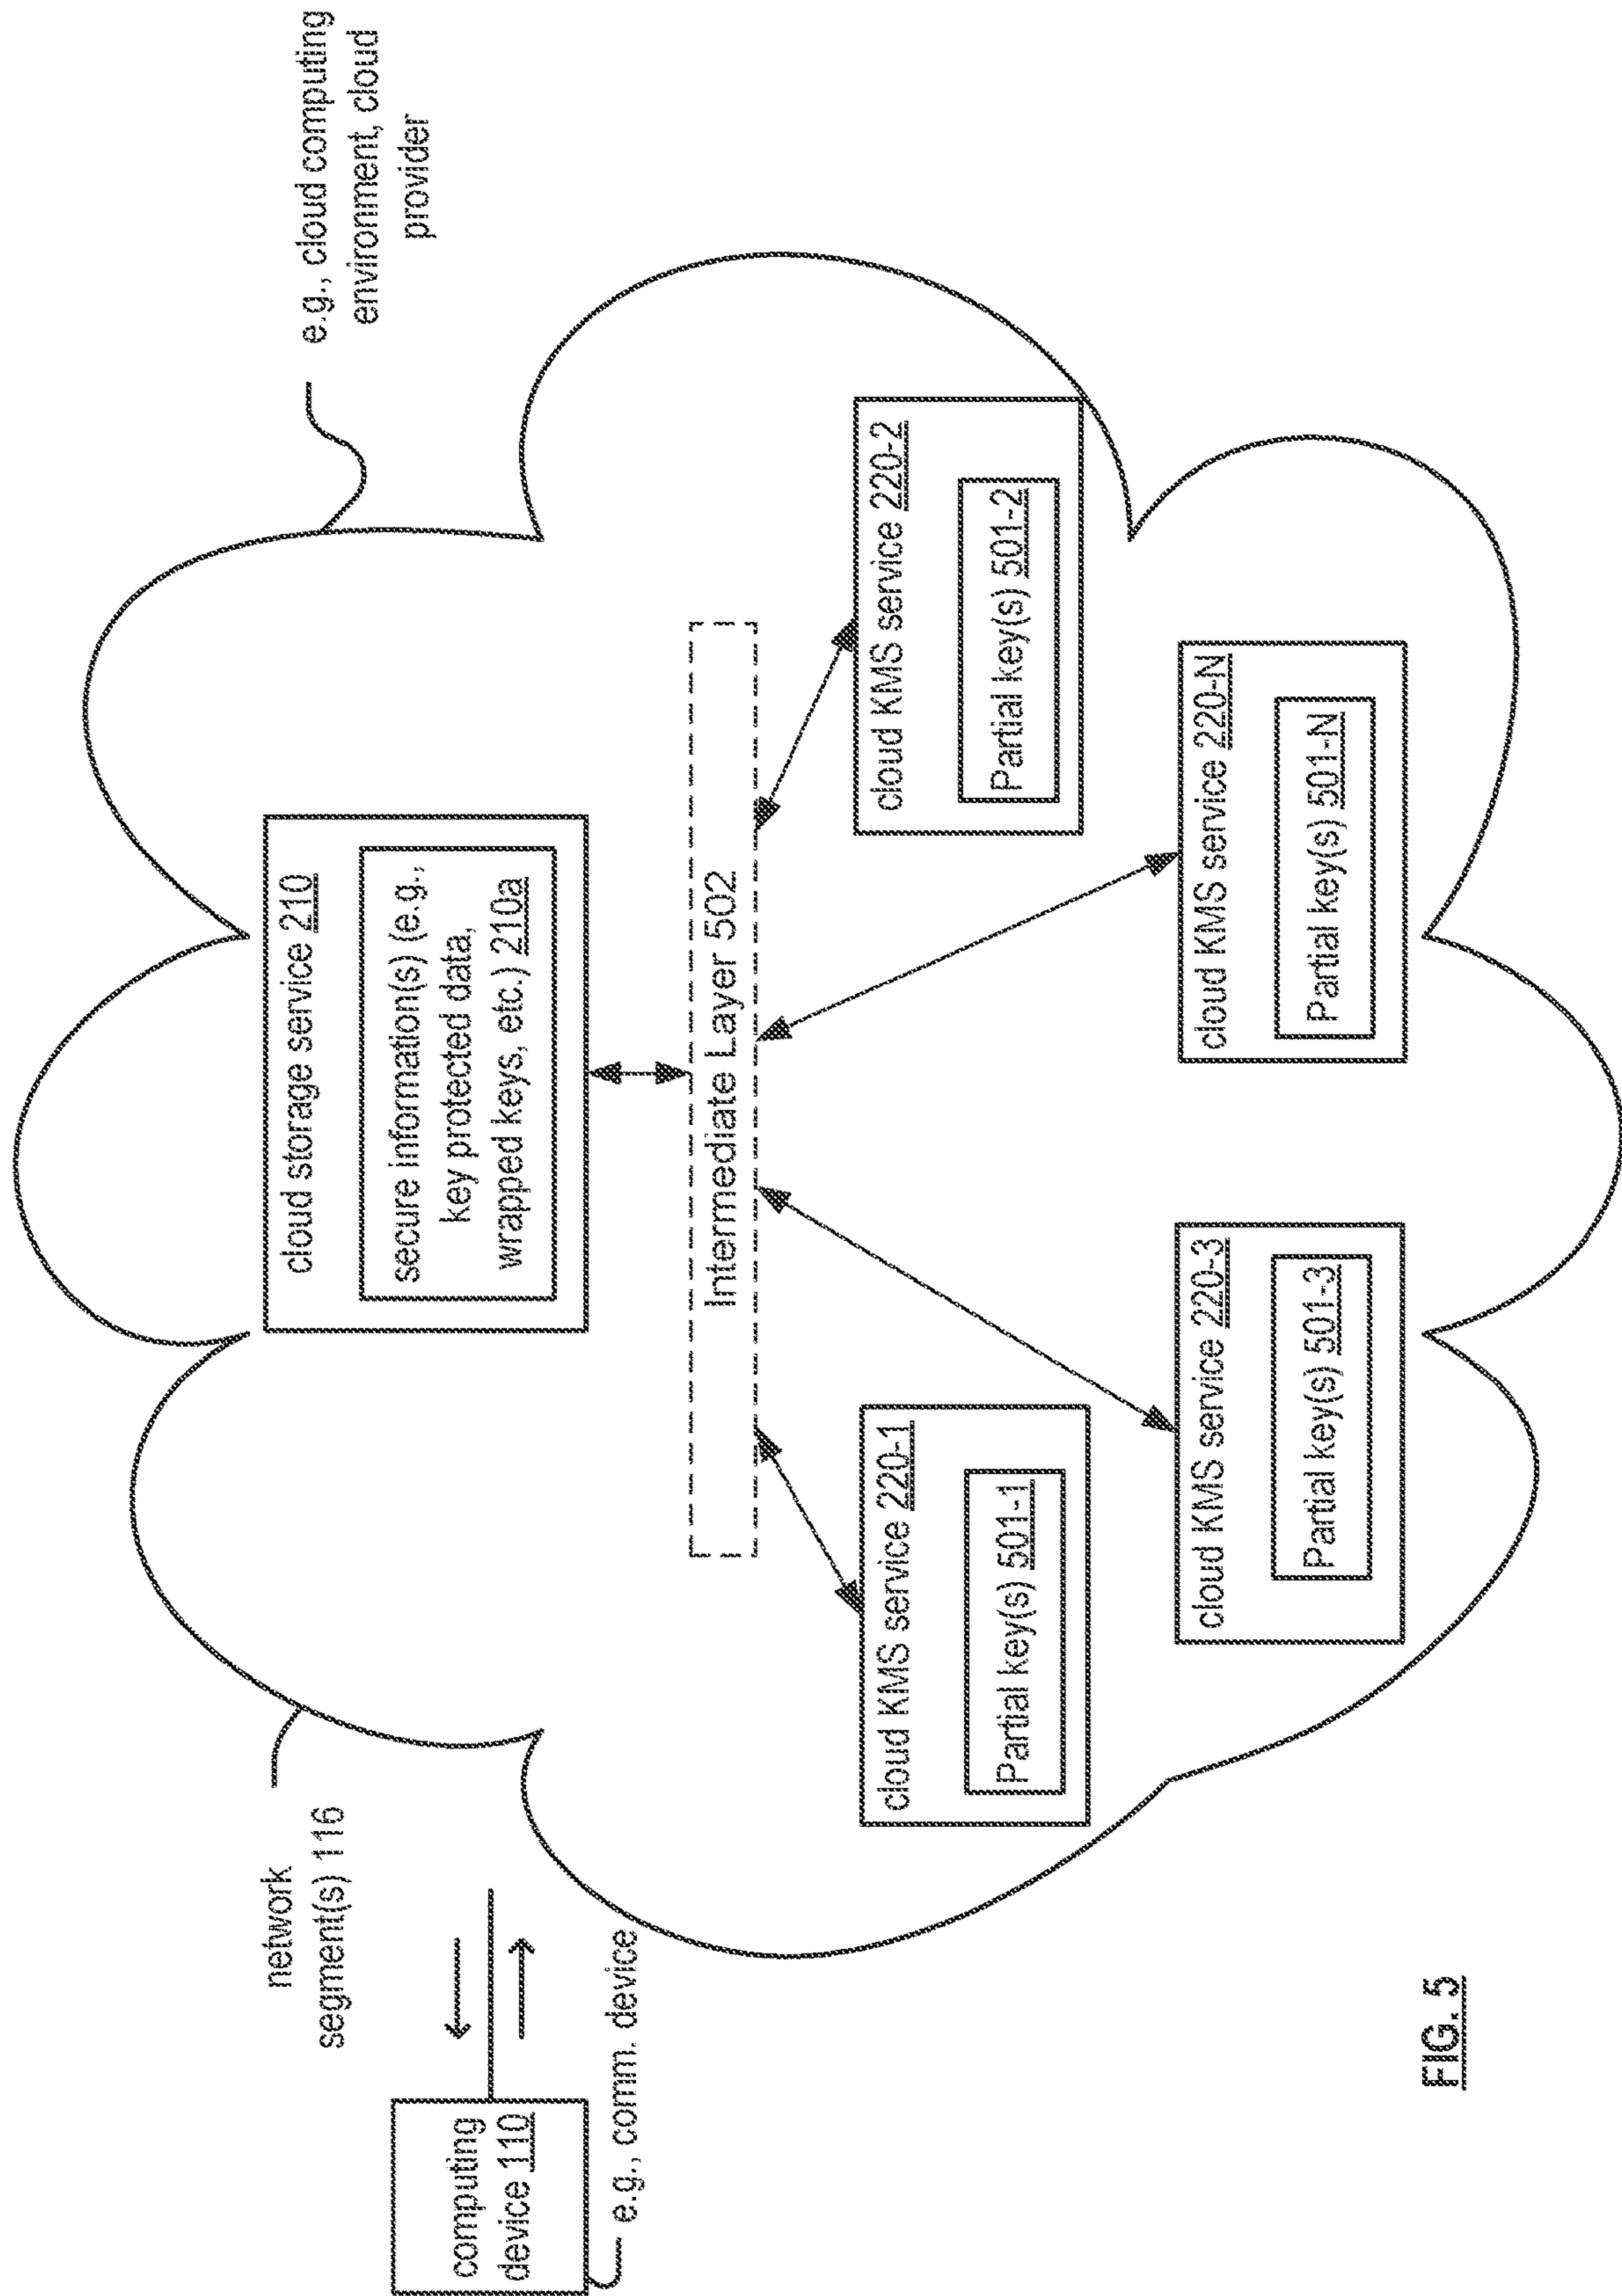


FIG. 5

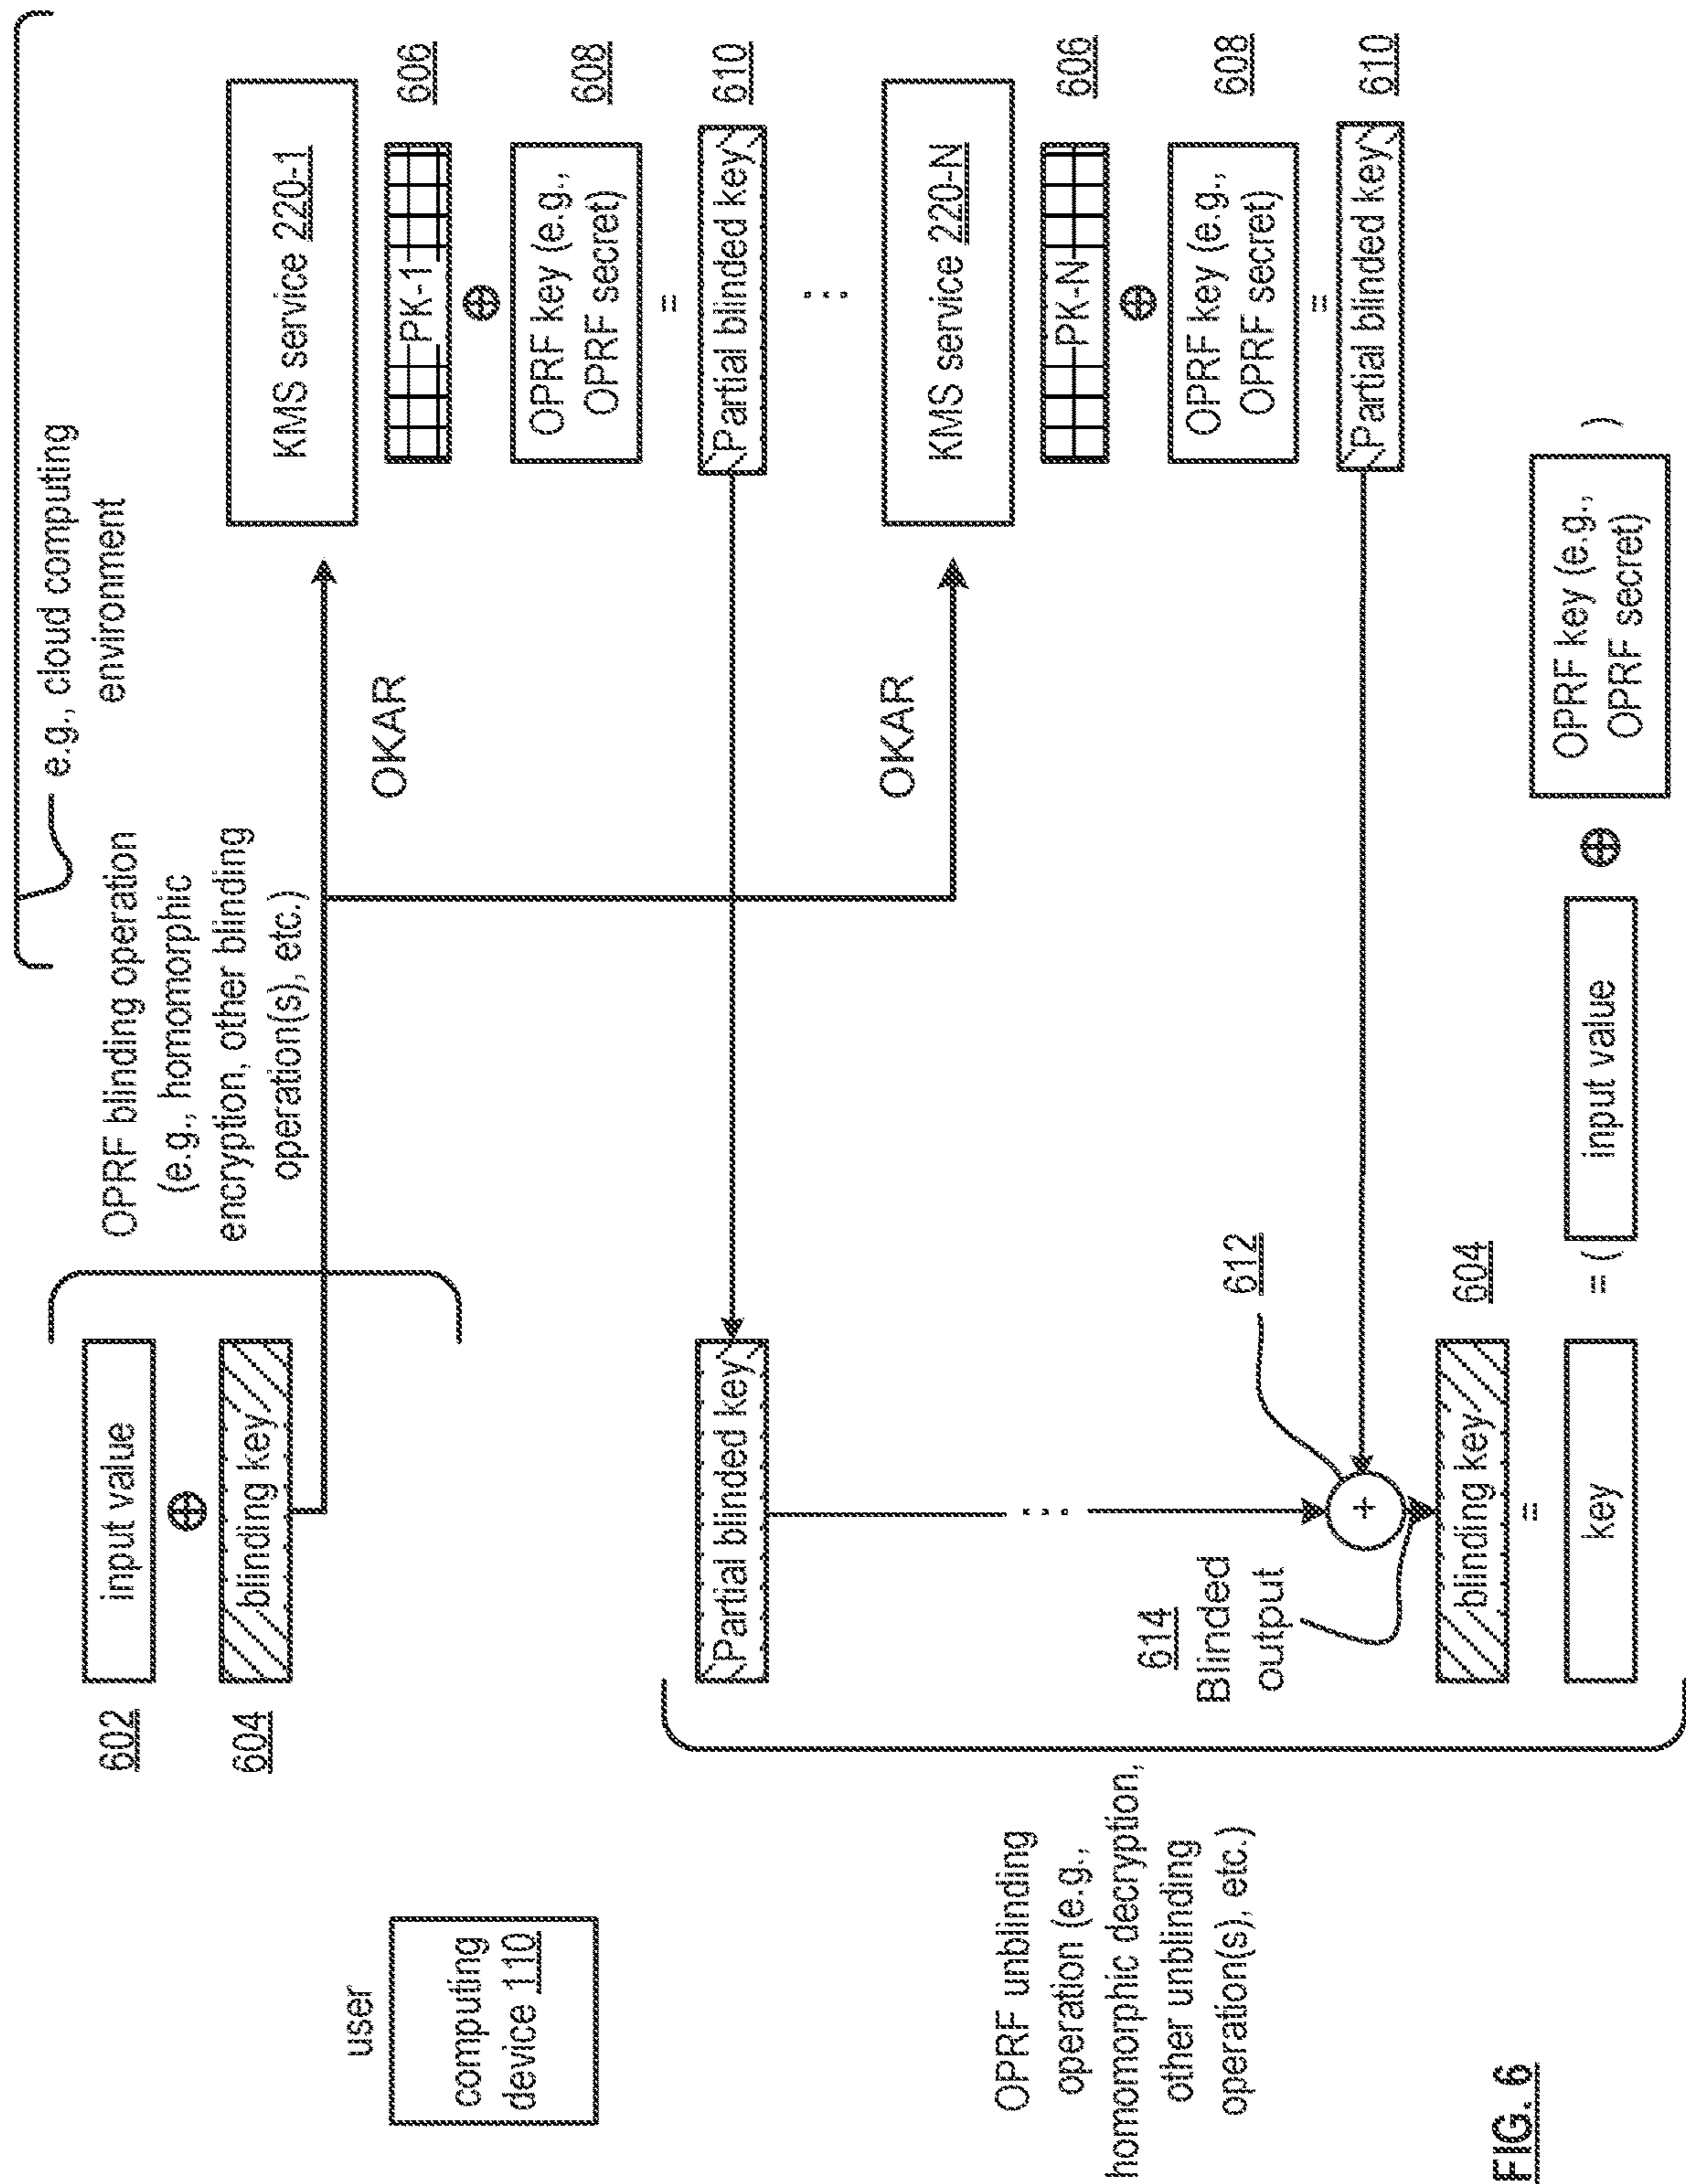


FIG. 6

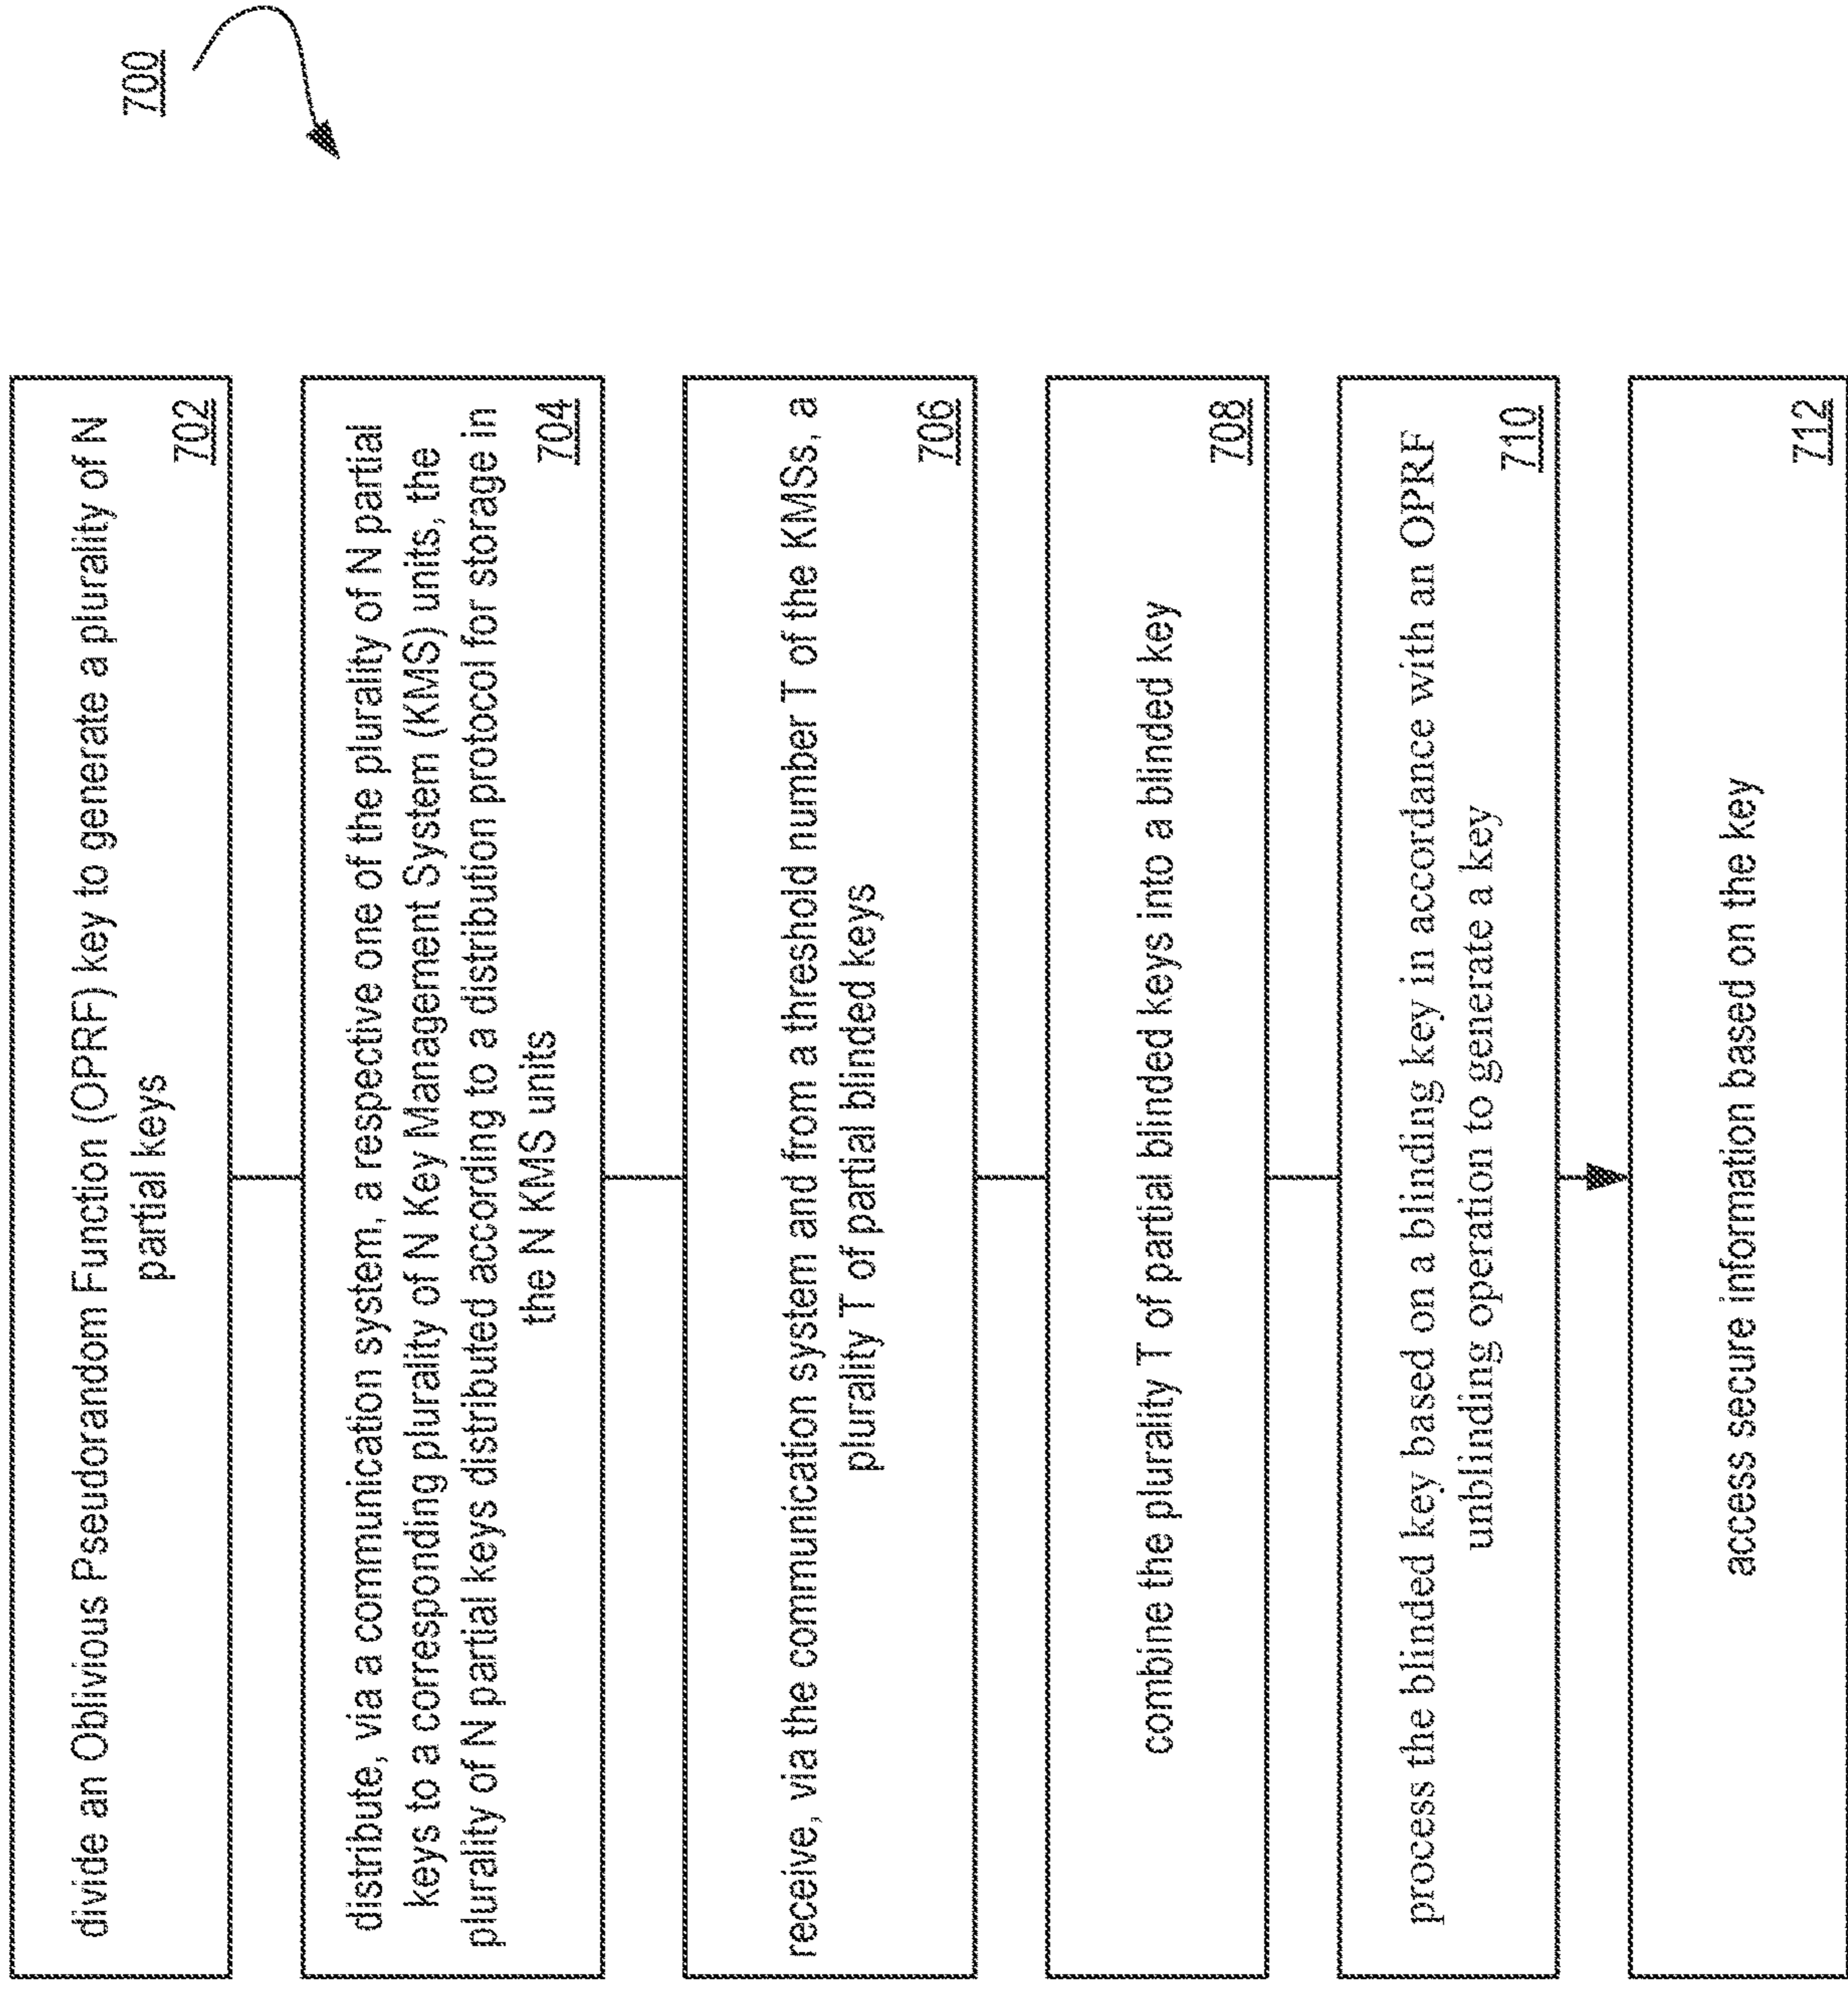


FIG. 7

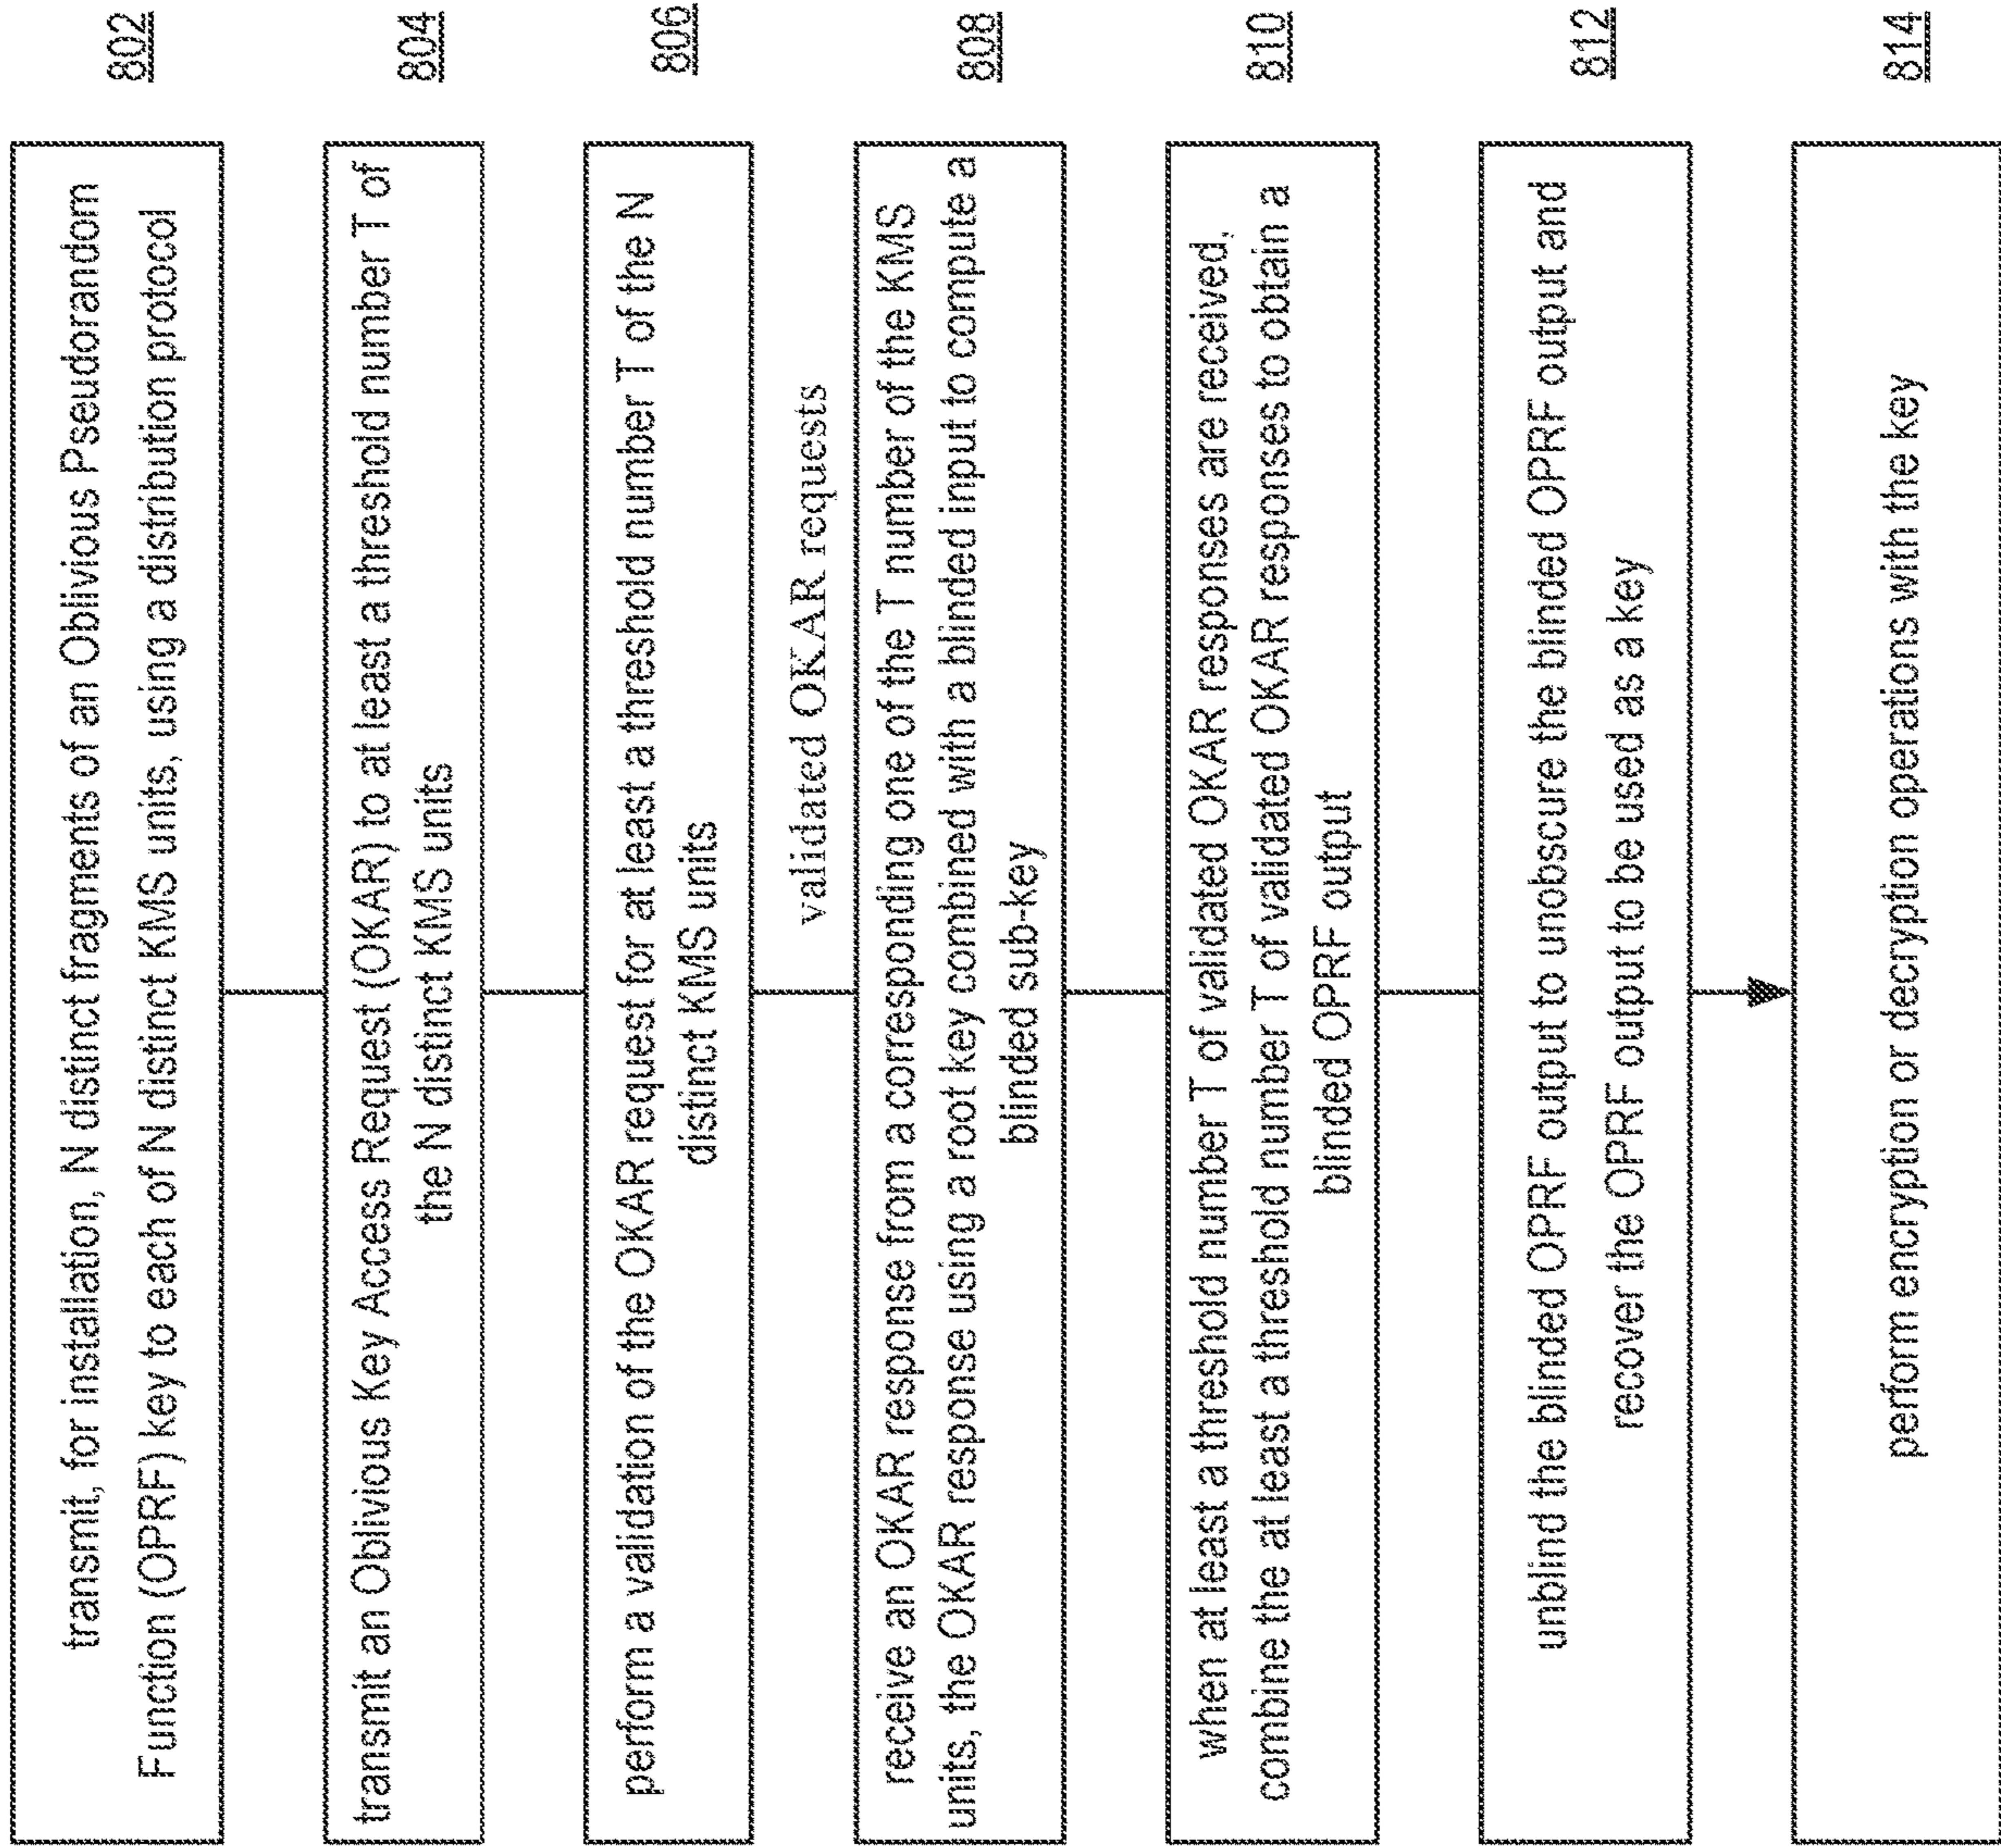


FIG. 8




FIG. 9




FIG. 10

1

THRESHOLD OBLIVIOUS PSEUDORANDOM FUNCTION IN A KEY MANAGEMENT SYSTEM

BACKGROUND

The technology described herein relates to security, encryption, and key management, and more specifically, to security, encryption, and key management in accordance with operations based on communication system and communications related to one or more Key Management Systems (KMSs).

In certain prior art communication system systems, ever-increasing quantities of data is stored online. Some data therein is critical, encrypted, secure, and/or private. For example, much of this data is private and some may be protected by confidentiality laws and regulations. Some of the data is encrypted to guard data from malicious insiders, external attackers, and/or accidental exposure requires. Encryption can operate using one or more encryption keys. Without appropriate encryption key(s), encrypted data cannot be deciphered. Therefore, reliable supply and management of keys is essential whenever dealing with encrypted data.

In addition, more recently, certain information is stored within one or more remote storage devices that are operated and maintained by another party. In certain prior art implementations, this other party service provider will have access to and be able to see the one or more encryption keys that is stores, manages, and provides to the clients and/or users that it services. In such situations, such a client and/or user can be totally susceptible and vulnerable to any bad intentions, behavior, lack of trust, etc. of such another party service provider.

Such prior art storage systems including those that store keys that are used to access secure and encrypted information (e.g., Key Management Systems (KMSs)) include many problems. In addition, within attempts to implement such storage systems based on cloud-based technology, the problems can become even more exacerbated. For example, cloud-based implementations of such prior art storage systems create additional security threats. Considering some examples, a tenant must trust the cloud provider and its operations. There must be trust that the cloud provider will manage the system correctly. This includes having faith that the cloud provider will be free of configuration errors, will effectively perform safeguard backups, will have correct policies and controls in place for tenant data, etc. Also, there must be confidence that the cloud provider will not have malicious insiders that may compromise sensitive data, will properly vet employees, will have auditing processes, and will not grant excessive unchecked power to system admins. Also, there must be assurance that the cloud provider will appropriately isolate tenants. This can include ensuring that the cloud provide will prevent side channel inter-virtual machine (VM) and inter-container attacks, protect memory, wipe drives between use, and prevent network sniffing. Also, there must be confidence that the cloud provider will guard against attacks including those that may originate from anywhere on Internet. There must be trust that the cloud provider will protect the network, (e.g., including the point of ingress for all cloud data), will monitor for breaches, will survive Distributed Denial of Service (DDoS), will apply patches, and will use and validate certificates appropriately.

Existing and prior art storage systems that are implemented based on cloud-based technology require placing significant trust in a single cloud provider and/or maintain-

2

ing one's own KMS infrastructure. Lack of security or confidence in adequate security are some of the major considerations that are preventing implementations of such prior art storage systems using on cloud-based technology. Such fears are existent when storing normal data (e.g., data that is not private, secret, and/or of high value), and those fears become amplified when dealing with secure data, encrypted data, encryption keys, etc. The prior art does not provide adequate solutions by which various concerns such as privacy, security, trust, etc. are adequately addressed while providing for confident and reliable user servicing for acceptable performance and user experience.

SUMMARY

Embodiments of the present invention disclose a computer-implemented method, a system, and a computer program product for effectuating one or more Key Management System (KMS) services based on one or more Threshold Oblivious Pseudorandom Functions (OPRFs). Within such an implementation, a plurality (threshold) of parties (e.g., multiple KMS units) operate cooperatively to evaluate a function and ultimately to provide a key to the client. A Threshold Oblivious Pseudorandom Function (T-OPRF) is an OPRF computed not by a single server, but by some quorum of servers, none of which has knowledge of the Key. In practical terms, this enables an OPRF system having no single point of failure and no single point of compromise. Computing the T-OPRF requires two steps, first the evaluation of the T-OPRF by each server in the quorum to produce a partial output, and second the combination of these partial outputs to yield the final output, which again is learned only by the client, regardless of who performs this combination. In any T-OPRF, there is the total number of servers "N" and the minimum quorum size "T", which defines a "T-of-N" threshold scheme.

In contradistinction, within prior art cloud-based technologies, the provider of a cloud-based KMS services operates such that the provider will see the tenant's keys. In the course of the computation, as performed using the novel solutions herein, without a threshold of KMSs, the various parties do not learn any information about an OPRF secret (e.g., an OPRF key). In certain optional operations, the OPRF secret is based on a Customer Root Key (CRK). Such novel solutions as presented herein including to support a Key Management System (KMS) can be implemented using cloud-based technology while requiring zero trust in the cloud provider and without requiring any KMS infrastructure in tenant premises.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1A is a diagram illustrating an embodiment of one or more communication systems supporting a Key Management System (KMS) according to various embodiments of the present invention;

FIG. 1B is a diagram illustrating an embodiment of one or more communication systems according to various embodiments of the present invention;

FIG. 1C is a diagram illustrating an embodiment of a computing device configured to operate within one or more communication systems according to various embodiments of the present invention;

FIG. 1D is a diagram illustrating an embodiment of a wireless communication system according to various embodiments of the present invention;

FIG. 2A is a diagram illustrating another embodiment of one or more communication systems supporting a KMS according to various embodiments of the present invention;

FIG. 2B is a diagram illustrating another embodiment of one or more communication systems supporting a KMS according to various embodiments of the present invention;

FIG. 3A is a diagram illustrating another embodiment of one or more communication systems supporting a KMS according to various embodiments of the present invention;

FIG. 3B is a diagram illustrating an embodiment of one or more communication systems supporting a KMS based on an Oblivious Pseudorandom Function (OPRF) according to various embodiments of the present invention;

FIG. 4A is a diagram illustrating an embodiment of one or more communication systems supporting key protect with obliviousness according to various embodiments of the present invention;

FIG. 4B is a diagram illustrating an embodiment of one or more communication systems supporting Hardware Security Module (HSM) integration according to various embodiments of the present invention;

FIG. 4C is a diagram illustrating an embodiment of key hierarchies as may be used in accordance with a KMS according to various embodiments of the present invention;

FIG. 5 is a schematic block diagram of a Threshold Oblivious Pseudorandom Function (T-OPRF) used in the process of producing encryption keys according to various embodiments of the present invention;

FIG. 6 is a diagram illustrating another embodiment of one or more communication systems supporting a multiple KMS unit structure based on a Threshold Oblivious Pseudorandom Function (T-OPRF) according to various embodiments of the present invention;

FIG. 7 is a diagram illustrating an embodiment of a method for execution by one or more computing devices according to various embodiments of the present invention;

FIG. 8 is a diagram illustrating another embodiment of a method for execution by one or more computing devices according to various embodiments of the present invention;

FIG. 9 depicts a cloud computing environment according to various embodiments of the present invention; and

FIG. 10 depicts abstraction model layers according to various embodiments of the present invention.

DETAILED DESCRIPTION

According to an embodiment of the present invention, novel key management and key security operates based on obliviousness to provide for security and secrecy of keys that are used such as based on one or more Key Management Systems (KMSs). This novel key management and key security also involves generation of one or more keys in accordance with an Oblivious Pseudorandom Function (OPRF) blinding operation (e.g., homomorphic encryption, one or more other blinding operations, etc.) and/or Oblivious Pseudorandom Function (OPRF) processing. In some examples, this includes performing an OPRF blinding operation (e.g., homomorphic encryption, one or more other blinding operations, etc.) (and an OPRF unblinding operation (e.g., homomorphic decryption, one or more other unblinding operations, etc.)) and supporting OPRF processing based on client (e.g., user) operations and service provider (e.g., KMS) operations. The various entities interact with one another in a manner that the two parties operate cooperatively to evaluate a function and ultimately to provide a key to the client. In accordance with such OPRF processing, the client cannot determine an OPRF secret used

by the service provider, and the service provider cannot determine the input and/or output associated with the client. In addition, the service provider cannot see and cannot determine the one or more keys that it provides to one or more clients. Also, in one example of such novel solutions as presented herein, the keys are regenerated each time they are used. This can be used to offload a large amount of storage and also to offload processing load to a large number of clients. This provides for the possible implementation of a scalable system that obviates the need for a centralized server or a Hardware Security Module (HSM) as is used in prior art systems.

FIG. 1A is a diagram illustrating an embodiment 100 of one or more communication systems supporting a Key Management System (KMS) according to various embodiments of the present invention. One or more computing devices (e.g., computing device 110, computing device 112, etc.) is configured to support communications via one or more other computing devices and/or one or more network segments 116. In one example, the computing device 110 is in communication with a computing device 112 via the one or more network segments 116. For example, the computing device 110 is associated with a first user, and the computing device 112 is associated with a second user. The one or more network segments 116 may be implemented in accordance with a cloud computing environment 50 such as described with reference to FIG. 10, and one or more network segments 116 may include one or more other computing devices therein (e.g., nodes, routers, gateways, servers, relays, transcoders, etc.) in some examples and/or embodiments.

The computing device 110 is configured to access secure information (e.g., secure, private, encrypted, etc. data, keys, etc.) based on one or more keys. Examples of such keys may be of various types including one or more of a Data Encryption Key (DEK), Key Encryption Key (KEK), Wrapped Data Encryption Key (WDEK), Master Key Encryption Key (M-KEK), Instance Key Encryption Key (I-KEK), Customer Root Key (CRK), and/or any other type of key including those associated with and used to encrypt and/or decrypt information, etc.

For example, once a key is generated, the computing device 110 may be configured to use that key to access secure information that is stored within the one or more network segments 116 and/or stored within a cloud-based technology that is based on or accessible via the one or more network segments 116. For example, the computing device 110 requests encrypted data that is stored by a cloud provider, receives that encrypted data that is stored by that cloud provider, and then uses the key to decrypt that encrypted data.

In general, in accordance with such security, encryption, etc., a key is used by the computing device 110 to access secure information (e.g., data, keys, etc.) that are kept unavailable to others that do not have the key. In general, a Key Management System (KMS) may be viewed as being a system for managing, reliably maintaining, and controlling access to keys on behalf computing devices, users, and/or applications, etc. High availability and durability is critical for a KMS. For example, considering a particular instance, if the KMS fails, any attempt to restore data encrypted with keys managed by the KMS will also fail. Security and proper access control enforcement and auditing is also essential. For example, if the wrong entity (e.g., an unauthorized entity) acquires a key from the KMS, the KMS has effectively disclosed to that party all data, keys, etc. encrypted under that key.

In an example of operation and implementation, a computing device **110** (e.g., such as associated with a user) is configured to support communications with a Key Management System (KMS) via one or more communication systems (e.g., the Internet, an intranet, a proprietary network, a private network, a public network, etc.) via the one or more network segments **116**. The computing device **110** is configured to process an input value that is associated with a key based on a blinding key in accordance with an Oblivious Pseudorandom Function (OPRF) blinding operation (e.g., homomorphic encryption, one or more other blinding operations, etc.) to generate a blinded value. Such an input value may be any desired value (e.g., an arbitrary identifier, a file name, an object name, an object identifier (ID), etc. a hash of any such elements such as a hash of an arbitrary identifier, a file name, an object name, an ID, etc.).

The computing device **110** is also configured to transmit (e.g., via the one or more network segments **116**) the blinded value to another computing device that is associated with a Key Management System (KMS) service (shown as KMS service **121**). The computing device **110** is also configured to receive (e.g., via the one or more network segments **116** and from the KMS service **121**) a blinded key. Note that the blinded key is based on processing of the blinded value based on an Oblivious Pseudorandom Function (OPRF) using an OPRF secret. In some examples, the KMS service **121** is configured to process the blinded value based on the OPRF using the OPRF secret. The computing device **110** is also configured to process the blinded key based on the blinding key in accordance with an OPRF unblinding operation (e.g., homomorphic decryption, one or more other unblinding operations, etc.) to generate the key that is associated with the input value. Then, the computing device **110** is also configured to access (e.g., via the one or more network segments **116**) secure information based on the key. Examples of such secure information may include information that is private, secret, and/or of high value. Examples of such data that may include any one or more of personal information, health data, financial data, strategic plans, competitive information, trade secrets, bitcoin wallets, customer lists, and/or any other type of data that may be private, secret, and/or of high value.

In some examples, the input value is unknown to the KMS service **121** and includes a key identifier (key id) that is associated with the key. Also, the key is unknown to the KMS service **121** and may include any type of key including those mentioned above including a DEK or a KEK. In some specific examples, the OPRF secret is unknown to the computing device **110** and is based on a Customer Root Key (CRK) that is associated with the computing device **110**.

FIG. 1B is a diagram illustrating an embodiment **102** of one or more communication systems according to various embodiments of the present invention. One or more network segments **116** provide communication inter-connectivity for at least two computing devices **110** and **112** (e.g., such computing devices may be implemented and operative to support communications with other computing devices in certain examples, and such computing devices may alternatively be referred to as communication devices in such situations including both computing device and communication device functionality and capability). Generally speaking, any desired number of communication devices are included within one or more communication systems (e.g., as shown by communication device **114**).

The various communication links within the one or more network segments **116** may be implemented using any of a variety of communication media including communication

links implemented as wireless, wired, optical, satellite, microwave, and/or any combination thereof, etc. communication links. In general, the one or more network segments **116** may be implemented to support a wireless communication system, a wire lined communication system, a non-public intranet system, a public internet system, a local area network (LAN), a wireless local area network (WLAN), a wide area network (WAN), a satellite communication system, a fiber-optic communication system, and/or a mobile communication system. Also, in some instances, communication links of different types may cooperatively form a connection pathway between any two communication devices. Considering one possible example, a communication pathway between devices **110** and **112** may include some segments of wired communication links and other segments of optical communication links. Note also that the devices **110-114** may be of a variety of types of devices including stationary devices, mobile devices, portable devices, etc. and may support communications for any of a number of services or service flows including data, telephony, television, Internet, media, synchronization, etc.

In an example of operation and implementation, device **110** includes a communication interface to support communications with one or more of the other devices **112-114**. In an example, the computing device **110** includes a communication interface configured to interface and communicate with a communication network (e.g., the one or more network segments **116**), memory that stores operational instructions, and a processor coupled to the communication interface and to the memory. The processor is configured to execute the operational instructions to perform various functions, operations, etc. Note that the communication supported by the computing device **110** may be bidirectional/to and from the one or more of the other computing devices **112-114** or unidirectional (or primarily unidirectional) from the one or more of the other computing devices **112-114**.

In one example, computing device **110** includes a processor that generates, modulates, encodes, etc. and transmits signals via a communication interface of the computing device **110** and also receives and processes, demodulates, decodes, etc. other signals received via the communication interface of the computing device **110** (e.g., received from other computing devices such as computing device **112**, computing device **114**, etc.).

Note also that the communication interface **120** may include functionality to support receipt of user input and output of user output (e.g., via a touchscreen, a keyboard, audio and/or video input/output ports, etc.) to facilitate interaction with one or more users of the computing device **110-1**. Such functionality is implemented separately from the communication interface **120** in some examples and is integrated into the communication interface **120** in other examples.

In an example of operation, computing device **110** is configured to process an input value that is associated with a key based on a blinding key in accordance with an OPRF blinding operation (e.g., homomorphic encryption, one or more other blinding operations, etc.) to generate a blinded value. The computing device **110** is also configured to transmit, via the communication system such as via the one or more network segments **116**, the blinded value to computing device **112** and/or computing device **114** that is associated with a Key Management System (KMS) service. The computing device **110** is also configured to receive, via the communication system such as via the one or more network segments **116** and from the computing device **112** and/or the computing device **114** that is associated with the

KMS service, a blinded key. Note that the blinded key is based on processing of the blinded value based on an Oblivious Pseudorandom Function (OPRF) using an OPRF secret. In some examples, the computing device **112** and/or computing device **114** is configured to process the blinded value based on the OPRF using the OPRF secret. Then, the computing device **110** is configured to process the blinded key based on the blinding key in accordance with an OPRF unblinding operation (e.g., homomorphic decryption, one or more other unblinding operations, etc.) to generate the key that is associated with the input value. In some examples, the computing device **110** is configured to use the key to access secure information (e.g., via the communication system such as via the one or more network segments **116**, based on locally available and/or stored secure information, and/or combination thereof, etc.). Such secure information may be data, keys, etc. and/or other information that is stored within one or more other computing devices that are accessible via the communication system such as via the one or more network segments **116**.

FIG. 1C is a diagram illustrating an embodiment **103** of a computing device configured to operate within one or more communication systems according to various embodiments of the present invention. The computing device **110-1** includes a communication interface **120** and processing circuitry **130**. The communication interface **120** includes functionality of a transmitter **122** and a receiver **124** to support communications with one or more other devices within a communication system. The computing device **110-1** may also include memory **140** to store information including one or more signals generated by the computing device **110-1** or such information received from other devices (e.g., computing device **112**) via one or more communication channels. For example, memory **140** may also include and store various operational instructions for use by the processing circuitry **130** in regard to the processing of messages and/or other received signals and generation of other messages and/or other signals including those described herein (e.g., image and/or video signals). Memory **140** may also store information including one or more types of encoding, one or more types of symbol mapping, concatenation of various modulation coding schemes, etc. as may be generated by the computing device **110-1** or such information received from other devices via one or more communication channels. The communication interface **120** supports communications to and from one or more other devices (e.g., computing device **112-1** and/or other computing devices). Memory **140** may also store information including one or more types of video and/or image processing in accordance with the various aspects, embodiments, and/or examples, and their equivalents, described herein.

Operation of the communication interface **120** may be directed by the processing circuitry **130** such that processing circuitry **130** transmits and receives signals (TX(s) and RX(s)) via the communication interface **120**. Generally speaking, computing device **110-1** is able to support communications with one or more other computing device within one or more communication systems including computing device **112-2**.

A computing device **110-1** (e.g., which may be any one of computing devices **110**, **112**, or **114** as with reference to FIG. 1B) is in communication with another computing device **112-1** (and/or any number of other wireless computing devices) via a communication medium. The computing device **110-1** includes a communication interface **120** to perform transmitting and receiving of at least one signal, symbol, packet, and/or frame, etc. (e.g., using a transmitter

122 and a receiver **124**) (note that general reference to packet or frame may be used interchangeably).

Generally speaking, the communication interface **120** is implemented to perform any such operations of an analog front end (AFE) and/or physical layer (PHY) transmitter, receiver, and/or transceiver. Examples of such operations may include any one or more of various operations including conversions between the frequency and analog or continuous time domains (e.g., such as the operations performed by a digital to analog converter (DAC) and/or an analog to digital converter (ADC)), gain adjustment including scaling, filtering (e.g., in either the digital or analog domains), frequency conversion (e.g., such as frequency upscaling and/or frequency downscaling, such as to a baseband frequency at which one or more of the components of the computing device **110-1** operates), equalization, pre-equalization, metric generation, symbol mapping and/or de-mapping, automatic gain control (AGC) operations, and/or any other operations that may be performed by an AFE and/or PHY component within a computing device.

In some implementations, the computing device **110-1** also includes a processing circuitry **130**, and an associated memory **140**, to execute various operations including interpreting at least one signal, symbol, packet, and/or frame transmitted to computing device **112-1** and/or received from the computing device **112-1** and/or any other computing device. The computing devices **110-1** and **112-1** may be implemented using at least one integrated circuit in accordance with any desired configuration or combination of components, modules, etc. within at least one integrated circuit. Also, the computing devices **110** and/or **112** may each include one or more antennas for transmitting and/or receiving of at least one packet or frame wirelessly (e.g., computing device **110-1** may include m antennas, and computing device **112-1** may include n antennas, where m and n are positive integers).

Also, in some examples, note that one or more of the processing circuitry **130**, the communication interface **120** (including the TX **122** and/or RX **124** thereof), and/or the memory **140** may be implemented in one or more “processing modules,” “processing circuits,” “processors,” and/or “processing units” or their equivalents. Considering one example, a system-on-a-chip (SOC) **130a** may be implemented to include the processing circuitry **130**, the communication interface **120** (including the TX **122** and/or RX **124** thereof), and the memory **140** (e.g., SOC **130a** being a multi-functional, multi-module integrated circuit that includes multiple components therein). Considering another example, processing-memory circuitry **130b** may be implemented to include functionality similar to both the processing circuitry **130** and the memory **140** yet the communication interface **120** is a separate circuitry (e.g., processing-memory circuitry **130b** is a single integrated circuit that performs functionality of a processing circuitry and a memory and is coupled to and also interacts with the communication interface **120**).

Considering even another example, two or more processing circuitries may be implemented to include the processing circuitry **130**, the communication interface **120** (including the TX **122** and/or RX **124** thereof), and the memory **140**. In such examples, such a “processing circuitry,” “processing circuitry,” or “processing circuitries” (or “processor” or “processors”) is/are configured to perform various operations, functions, communications, etc. as described herein. In general, the various elements, components, etc. shown within the computing device **110-1** may be implemented in any number of “processing modules,” “processing circuits,”

“processors,” and/or “processing units” (e.g., 1, 2, . . . , and generally using N such “processing modules,” “processing circuits,” “processors,” and/or “processing units”, where N is a positive integer greater than or equal to 1).

In some examples, the computing device **110-1** includes both processing circuitry **130** and communication interface **120** configured to perform various operations. In other examples, the computing device **110-1** includes SOC **130a** configured to perform various operations. In even other examples, the computing device **110-1** includes processing-memory circuitry **130b** configured to perform various operations. Generally, such operations include generating, transmitting, etc. signals intended for one or more other computing device (e.g., computing device **112-1**) and receiving, processing, etc. other signals received for one or more other devices (e.g., computing device **112-1**).

In some examples, note that the communication interface **120**, which is coupled to the processing circuitry **130**, is configured to support communications within a satellite communication system, a wireless communication system, a wired communication system, a fiber-optic communication system, and/or a mobile communication system (and/or any other type of communication system implemented using any type of communication medium or media). Any of the signals generated and transmitted and/or received and processed by the computing device **110-1** may be communicated via any of these types of communication systems.

Note that computing device **110-1** may be implemented to operate as any one or more of a satellite communication device, a wireless communication device, a wired communication device, a fiber-optic communication device, or a mobile communication device and implemented and/or operative within any one or more communication systems including a satellite communication system, a wireless communication system, a wired communication system, a fiber-optic communication system, or a mobile communication system, among other types of communication systems.

In an example of operation and implementation, a computing device includes a communication interface **120** configured to interface and communicate with a communication network, memory **140** that stores operational instructions, and processing circuitry **130** coupled to the communication interface and to the memory.

The processing circuitry **130** is configured to execute the operational instructions to perform various functions, operations, and processes (sometimes in cooperation with the communication interface **120** and/or the memory **140**). In an example, the processing circuitry **130** is configured to process an input value that is associated with a key based on a blinding key in accordance with an OPRF blinding operation (e.g., homomorphic encryption, one or more other blinding operations, etc.) to generate a blinded value. The processing circuitry **130** is then configured to transmit, via the communication interface **120**, the blinded value to another computing device **112-1** that is associated with a Key Management System (KMS) service. The processing circuitry **130** is configured to receive, via the communication interface **120** and from the computing device **112-1** that is associated with the KMS service, a blinded key. Note that the blinded key is based on processing of the blinded value based on an OPRF using an OPRF secret. In some examples, the computing device **112-1** is configured to process of the blinded value based on an OPRF using an OPRF secret. The processing circuitry **130** is then configured to process the blinded key based on the blinding key in accordance with an OPRF unblinding operation (e.g., homomorphic decryption, one or more other unblinding operations, etc.) to generate the key

that is associated with the input value. In some examples, the computing device **110-1** is configured to access secure information based on the key (e.g., via the communication system, based on locally available and/or stored secure information, and/or combination thereof, etc.).

FIG. 1D is a diagram illustrating an embodiment **100** of a wireless communication system according to various embodiments of the present invention. The wireless communication system includes one or more base stations and/or access points **150**, wireless communication devices **160-166** (e.g., wireless stations (STAs)), and a network hardware component **156**. The wireless communication devices **160-166** may be laptop computers, or tablets, **160**, personal digital assistants **162**, personal computers **164** and/or cellular telephones **166** (and/or any other type of wireless communication device). Other examples of such wireless communication devices **160-166** could also or alternatively include other types of devices that include wireless communication capability (and/or other types of communication functionality such as wired communication functionality, satellite communication functionality, fiber-optic communication functionality, etc.). Examples of wireless communication devices may include a wireless smart phone, a cellular phone, a laptop, a personal digital assistant, a tablet, a personal computers (PC), a work station, and/or a video game device.

Some examples of possible devices that may be implemented to operate in accordance with any of the various examples, embodiments, options, and/or their equivalents, etc. described herein may include, but are not limited by, appliances within homes, businesses, etc. such as refrigerators, microwaves, heaters, heating systems, air conditioners, air conditioning systems, lighting control systems, and/or any other types of appliances, etc.; meters such as for natural gas service, electrical service, water service, Internet service, cable and/or satellite television service, and/or any other types of metering purposes, etc.; devices wearable on a user or person including watches, monitors such as those that monitor activity level, bodily functions such as heart-beat, breathing, bodily activity, bodily motion or lack thereof, etc.; medical devices including intravenous (IV) medicine delivery monitoring and/or controlling devices, blood monitoring devices (e.g., glucose monitoring devices) and/or any other types of medical devices, etc.; premises monitoring devices such as movement detection/monitoring devices, door closed/ajar detection/monitoring devices, security/alarm system monitoring devices, and/or any other type of premises monitoring devices; multimedia devices including televisions, computers, audio playback devices, video playback devices, and/or any other type of multimedia devices, etc.; and/or generally any other type(s) of device(s) that include(s) wireless communication capability, functionality, circuitry, etc. In general, any device that is implemented to support wireless communications may be implemented to operate in accordance with any of the various examples, embodiments, options, and/or their equivalents, etc. described herein.

The one or more base stations (BSs) or access points (APs) **150** are operably coupled to the network hardware **156** via local area network connection **152**. The network hardware **156**, which may be a router, switch, bridge, modem, system controller, etc., provides a wide area network connection **154** for the communication system. Each of the one or more base stations or access points **150** has an associated antenna or antenna array to communicate with the wireless communication devices in its area. Typically, the wireless communication devices register with a particular

base station or access point **150** to receive services from the communication system. For direct connections (i.e., point-to-point communications), wireless communication devices communicate directly via an allocated channel.

Any of the various wireless communication devices (WDEVs) **160-166** and one or more BSs or APs **150** may include a processing circuitry and/or a communication interface to support communications with any other of the wireless communication devices **160-166** and one or more BSs or APs **150**. In an example of operation, a processing circuitry and/or a communication interface implemented within one of the devices (e.g., any one of the WDEVs **160-166** and one or more BSs or APs **150**) is/are configured to process at least one signal received from and/or to generate at least one signal to be transmitted to another one of the devices (e.g., any other one of the one or more WDEVs **160-166** and one or more BSs or APs **150**).

Note that general reference to a communication device, such as a wireless communication device (e.g., WDEVs) **160-166** and one or more BSs or APs **150** in FIG. 1D, or any other communication devices and/or wireless communication devices may alternatively be made generally herein using the term 'device' (e.g., "device" when referring to "wireless communication device" or "WDEV"). Generally, such general references or designations of devices may be used interchangeably.

The processing circuitry and/or the communication interface of any one of the various devices, WDEVs **160-166** and one or more BSs or APs **150**, may be configured to support communications with any other of the various devices, WDEVs **160-166** and one or more BSs or APs **150**. Such communications may be uni-directional or bi-directional between devices. Also, such communications may be uni-directional between devices at one time and bi-directional between those devices at another time.

In an example, a device (e.g., any one of the WDEVs **160-166** and one or more BSs or APs **150**) includes a communication interface and/or a processing circuitry (and possibly other possible circuitries, components, elements, etc.) to support communications with other device(s) and to generate and process signals for such communications. The communication interface and/or the processing circuitry operate to perform various operations and functions to effectuate such communications (e.g., the communication interface and the processing circuitry may be configured to perform certain operation(s) in conjunction with one another, cooperatively, dependently with one another, etc. and other operation(s) separately, independently from one another, etc.). In some examples, such a processing circuitry includes all capability, functionality, and/or circuitry, etc. to perform such operations as described herein. In some other examples, such a communication interface includes all capability, functionality, and/or circuitry, etc. to perform such operations as described herein. In even other examples, such a processing circuitry and a communication interface include all capability, functionality, and/or circuitry, etc. to perform such operations as described herein, at least in part, cooperatively with one another.

In an example of implementation and operation, a wireless communication device (e.g., any one of the WDEVs **160-166** and one or more BSs or APs **150**) includes a processing circuitry to support communications with one or more of the other wireless communication devices (e.g., any other of the WDEVs **160-166** and one or more BSs or APs **150**). For example, such a processing circuitry is configured to perform both processing operations as well as commu-

nication interface related functionality. Such a processing circuitry may be implemented as a single integrated circuit, a system on a chip, etc.

In another example of implementation and operation, a wireless communication device (e.g., any one of the WDEVs **160-166** and one or more BSs or APs **150**) includes a processing circuitry, a communication interface, and a memory configured to support communications with one or more of the other wireless communication devices (e.g., any other of the WDEVs **160-166** and one or more BSs or APs **150**).

In an example of operation and implementation, WDEV **160** is configured to process an input value that is associated with a key based on a blinding key in accordance with an OPRF blinding operation (e.g., homomorphic encryption, one or more other blinding operations, etc.) to generate a blinded value. The WDEV **160** is then configured to transmit the blinded value to another computing device (e.g., WDEV **164** and/or some other computing device via network hardware **156**) that is associated with a Key Management System (KMS) service. The WDEV **160** is configured to receive, from that other computing device that is associated with the KMS service, a blinded key. Note that the blinded key is based on processing of the blinded value based on an OPRF using an OPRF secret. In some examples, this computing device is configured to process of the blinded value based on an OPRF using an OPRF secret. The WDEV **160** is then configured to process the blinded key based on the blinding key in accordance with an OPRF unblinding operation (e.g., homomorphic decryption, one or more other unblinding operations, etc.) to generate the key that is associated with the input value. In some examples, the WDEV **160** is configured to access secure information based on the key (e.g., via the communication system, based on locally available and/or stored secure information, and/or combination thereof, etc.).

In accordance with a Key Management System (KMS), and particularly in accordance with a KMS that operates based on a cloud-based system, there can be situations in which there is fear, uncertainty, lack of confidence, and/or lack of trust, etc. when storing normal data in the cloud are amplified and particularly when it comes to storage of the encryption keys that are used to access that data. For example, some data that may be encrypted may include any one or more of personal information, health data, financial data, strategic plans, competitive information, trade secrets, bitcoin wallets, customer lists, and/or other types of data. Such data is encrypted when it is private, secret, or of high-value. as such, the security of encrypted data is directly related to and dependent on, at least in part, on the security of the keys used to access that data.

FIG. 2A is a diagram illustrating another embodiment **201** of one or more communication systems supporting a KMS according to various embodiments of the present invention. This diagram shows a computing device **110** that is configured to interact with a cloud storage service **210** and a cloud Key Management System (KMS) service **220** that are both implemented within the same environment (e.g., one or more network segments **116** that may be implemented as a cloud computing environment, a cloud provider, etc.). The cloud storage service **210** may include various types of one or more secure information **210a** (e.g., key protected data, wrapped key, etc. and/or other secure information). The cloud KMS service **220** may include one or more keys **220a** that may be used by one of more users associated with one or more computing devices to access the various types of one or more secure information **210a**.

In this implementation, when the same cloud provider is used for both the cloud storage service **210** and the cloud KMS service **220**, a malicious insider or corrupt cloud provider could access tenant data. As such, trust requirements are maximized in order to ensure the security of the data. For example, a customer using the same provider for both services (e.g., the cloud storage service **210** and the cloud KMS service **220**) requires that they have complete confidence in that provider and its processes.

FIG. 2B is a diagram illustrating another embodiment **202** of one or more communication systems supporting a KMS according to various embodiments of the present invention. This diagram shows a computing device **110** that is configured to interact with a cloud storage service **210** and a cloud KMS service **220** that are separately implemented within different environments (e.g., the cloud storage service **210** implemented based on one or more network segments **116** that may be implemented as a first cloud computing environment, a cloud provider, etc., and the KMS service **220** based on one or more network segments **116a** that may be implemented as a second cloud computing environment, a cloud provider, etc.). Similarly as described above, the cloud storage service **210** may include various types of one or more secure information **210a** (e.g., key protected data, wrapped key, etc. and/or other secure information), and the cloud KMS service **220** may include one or more keys **220a** that may be used by one of more users associated with one or more computing devices to access the various types of one or more secure information **210a**.

In this implementation, when two separate and distinct cloud providers are used for the cloud storage service **210** and a cloud KMS service **220**, respectively, there can be higher complexity of the overall system, and there can be a likelihood of incompatible Application Program Interfaces (APIs). As such certain interoperability issues and other problems may emerge. For example, note that while customer doesn't necessarily need to have complete trust in any one specific provider, this implementation can unfortunately introduce a number of interoperability issues. For example, APIs and libraries for interfacing between the two services may not be compatible. Also, functions such as "Server-Side Encryption" (SSE) may not work at all in such an implementation.

FIG. 3A is a diagram illustrating another embodiment **301** of one or more communication systems supporting a KMS according to various embodiments of the present invention. This diagram shows a computing device **110** that is configured to interact with a cloud storage service **310** and an internal KMS service **320a** that may be implemented at a tenant premises **330**. For example, the tenant premises **330** may be located remotely from the computing device **110** and is accessible via one or more network segments **116** that may be implemented as a cloud computing environment, a cloud provider, etc. Similarly as described above with respect to other embodiments and examples, the cloud storage service **310** may include various types of one or more secure information **310a** (e.g., key protected data, wrapped key, etc. and/or other secure information), and the internal KMS service **320** may include one or more keys **320b** that may be used by one of more users associated with one or more computing devices to access the various types of one or more secure information **310a**.

In this implementation, when a tenant uses cloud storage (e.g., cloud storage service **310**) to operate with an internal KMS service **320**, the customer needs not necessarily have complete trust with the cloud provider, but such an imple-

mentation can require significant processes, expertise and expense to manage one's own KMS. Such an implementation can be very expensive.

For example, this diagram shows an implementation of a user of the cloud that may have little to no trust in the cloud provider to protect the security of its one or more keys. For example, such a user may trust a cloud storage provider to store encrypted data, yet that user trusts no cloud provider with holding of its one or more keys. In this implementation, while the user doesn't necessarily have to trust the cloud provider, the implementation can be problematic for various reasons including being very expensive, requiring a rare expertise, special-purpose equipment, requiring disaster recovery plan(s), trusted staff, and rigorous policies. Without these, it is very likely to be less reliable or less secure in practice than a cloud KMS.

Many such implementations of KMSs based on cloud-based technologies suffer for various reasons including requiring placing significant trust in a single provider and/or requiring the maintenance of one's own KMS infrastructure. This disclosure addresses such deficiencies and problems in the prior art including to provide a KMS that does not require any trust in the cloud provider and also needs no KMS infrastructure in the tenant premises. Such novel solutions as presented herein minimizes any required trust in a KMS service provider. For example, the one or more keys never leave customer premises, and the KMS service provider never sees those one or more keys. Also, no one can access the one or more keys without authenticating to the KMS provider. In addition, such novel solutions as presented herein provides for post-quantum security, in that, even with the advent of quantum capability of performing near limitless computation operations, the novel implementation as presented herein is immune to such advances in computing technology as may be targeted for hacking, invasive processes, etc. For example, a novel key access protocol approach as presented herein is immune from attackers with unlimited computational resources, including those with quantum computers.

Also, within such novel solutions as presented herein, the security of the keys if not dependent on the security of one or more communication channels over which communications are made. For example, some prior art approaches operate based on Transport Layer Security (TLS) and/or other means by to effectuate secure communications.

In addition, such novel solutions as presented herein provides for everlasting security, in that, the one or more keys remain secure. For example, even in the unfortunate event in which a KMS service provider is completely breached, the one or more keys remain totally secure. Note that some implementations may be implemented as requiring unpredictable key identifiers (ids) (e.g., using a "key id as a second factor"). Such novel solutions as presented herein obviates the requirement to have full and complete trust in a KMS service provider in terms of using or exposing the one of more keys the user wants the KMS service provider to store. The user may still seek a KMS service provider that does provide a highly available/reliable system, but the trust in that same KMS service provider to trust fully the KMS service provider in terms of using or exposing the one of more keys is obviated.

In an implementation when the tenant of a cloud KMS service provider trusts the provider with his keys, the tenant may operate by either storing them for later retrieval or in unwrapping them. In both cases, the cloud KMS service provider will see the tenant's keys. Such novel solutions as

presented herein provides for a means by which the cloud KMS service provider will never see the tenant's keys.

FIG. 3B is a diagram illustrating an embodiment 302 of one or more communication systems supporting a KMS based on an Oblivious Pseudorandom Function (OPRF) according to various embodiments of the present invention. An OPRF enables the tenant to get keys from the Cloud KMS provider. The property of obliviousness ensures the provider is cryptographically, mathematically, and provably not able to see or determine the keys. With respect to this diagram note that the "+" and "-" operations depicted therein are not arithmetical addition and subtraction per se. These operations may be exponentiation modulo a prime, or multiplication over an elliptic curve, or some other operations.

Considering an OPRF, an OPRF allows two parties to evaluate a function, Y, as follows:

$$Y = \text{OPRF}(K, X)$$

The OPRF secret K is only known to "Bob"; Alice can't determine it.

Output Y and input X are only known to "Alice"; Bob can't determine either.

An Oblivious PRF enables an ideal Cloud KMS:

The tenant uses "X" as a "key id" and "Y" as the key (DEK or KEK)

The OPRF guarantees the provider learns nothing about the key

The provider holds the OPRF secret: "K", functioning as a "CRK"

The OPRF can be viewed a key derivation that occurs on a blinded value (e.g., on a homomorphically encrypted cipher text). Note that certain examples herein are described with respect to an OPRF blinding operation that is performed using homomorphic encryption (and a an OPRF unblinding operation that is performed using homomorphic decryption), in general, any OPRF blinding/OPRF unblinding operation may be used such that the process that performs the OPRF blinding/OPRF unblinding operation are known to client (e.g., user, computing device 110 such as "Alice," and not to "Bob"). One example of such OPRF blinding/OPRF unblinding operation includes homomorphic encryption/homomorphic decryption. However, in general, any function may be used by the perform the OPRF blinding/OPRF unblinding operation to generate a blinded value that is unknown to the other computing device that is associated with the KMS service (e.g., server/KMS service 121 ("Bob")). For example, any desired function or mapping of an input value to generate an unknown input value (unknown to the other computing device that is associated with the KMS service such as server/KMS service 121 ("Bob")). Then, the client (e.g., user, computing device 110) knows how to perform the appropriate OPRF unblinding based on the OPRF blinding that was performed in that instance.

With respect to such an OPRF blinding operation that is performed using homomorphic encryption, the client (e.g., user, computing device 110) applies a homomorphic one-time-pad encryption key to an input value. For example, the client starts with some input value from which it wants to derive a key. For example, this input value could be a key id. The client then encrypts the input value with a one-time-pad encryption key (e.g., homomorphic one-time-pad encryption). The one-time-pad encryption key is randomly generated for this key recovery only and is or may be thrown away afterwards (e.g., not saved for future use).

This is an OPRF blinding operation that is performed using homomorphic encryption, and a one-time-pad. Accordingly, the encrypted (or blinded) result reveals zero information about the input. This cipher text (blinded value) is then sent to the server/KMS service 121. The server/KMS service 121 uses the OPRF key (e.g., a Customer Root Key (CRK) in some examples, an OPRF secret) to perform a key derivation function on this cipher text and returns it to the user.

Because of the homomorphic properties of the encryption, when the client (e.g., user, computing device 110) decrypts the result from the server/KMS service 121, it finds it gets the same value as had the server performed its key derivation function directly against the plaintext value. This result is considered the key. The resulting key is equal to the key derivation function (KDF) (e.g., a deterministic function used to generate the key) applied to the input.

Note that even if the client had chosen a different random blinding key, note that the resulting key that would be generated would be the same. In fact, all possible blinded values are possible with any possible input value. This is why the server, and any eavesdroppers, gain no information about the input value or the derived key, from seeing what goes over the wire, network segment(s), cloud, etc. to the server/KMS service 121.

Note that this form of homomorphic encryption is a special-case form that is extremely efficient and practical. Note also that Hardware Security Modules (HSMs) are quoted as capable of performing 10s of thousands of such operations per second. CPUs can perform upwards of hundreds of thousands per second.

Also, note that the Key is derived from the "Input Value" combined with the "OPRF Key" (e.g., an OPRF secret). Note also that the holder of the OPRF Key, the server/KMS service 121, never sees the Input Value, nor the Resulting key. This is enforced by a process of "Blinding" where both the input and the output are blinded in an information theoretically (quantum secure) way, such that the input and output yields zero information about the Key, neither to the KMS service, hackers, or the NSA. The following steps may be viewed as effectuating this process and exchange between a client (e.g., user, computing device 110) and server/KMS service 121:

1. Tenant generates random key: R
2. Tenant encrypts the "key id" using the random key: $E_R\{\text{key-id}\}$
3. Tenant sends encrypted result to the Cloud KMS provider
4. Cloud KMS provider encrypts result with its own key: P
5. Cloud KMS provider returns the result to the tenant: $E_P\{E_R\{\text{key-id}\}\}$
6. Tenant decrypts it with his random key R to get: $E_P\{\text{key-id}\}$

In some examples, a computing device 110 (e.g., a client such as associated with a user) is configured to process an input value that is associated with a key based on a blinding key in accordance with homomorphic encryption to generate a blinded value. The computing device is configured to transmit, via a communication system, the blinded value to another computing device (server/KMS service 121) that is associated with a Key Management System (KMS) service. The computing device is configured to receive, via the communication system and from the other computing device (server/KMS service 121) that is associated with the KMS service, a blinded key. The blinded key is based on processing of the blinded value based on an OPRF using an OPRF

secret. For example, the server/KMS service **121** is configured to on processing of the blinded value based on the OPRF using the OPRF secret. The computing device **110** is then configured to process the blinded key based on the blinding key in accordance with homomorphic decryption to generate the key that is associated with the input value. In some examples, the computing device **110** is also configured to access, via the communication system, secure information based on the key.

In an example of operation and implementation, once the key is generated, the computing device **110** uses that key to access secure information that is stored within a cloud-based technology that is based on or accessible via the one or more network segments. For example, the computing device **110** requests encrypted data that is stored by a cloud provider, receives that encrypted data that is stored by that cloud provider, and then uses the key to decrypt that encrypted data.

Further understanding of an Oblivious Pseudorandom Function (OPRF) may be made based on consideration of a Pseudorandom Function (PRF) (e.g., that is not oblivious). A Pseudorandom Function (PRF) is a function that takes two inputs:

1. a PRF key “K”; and
2. an arbitrary message “M”.

From these two inputs, the PRF returns a “pseudorandom” output. This is an output that is statistically indistinguishable from random output. Also, the output is infeasible to be predicted without knowledge of K. These two properties make PRFs well-suited for key derivation, that is, creating sub-keys from some top-level “root” key. For example, an unlimited number of sub-keys may be computed from a PRF as follows:

$$\text{sub-key}_1 = \text{PRF}(K, "1"), \text{sub-key}_2 = \text{PRF}(K, "2"), \text{sub-key}_3 = \text{PRF}(K, "3"), \dots, \text{sub-key}_n = \text{PRF}(K, "n")$$

This can simplify key management, as only a single top-level, or root key needs to be persisted while supporting a virtually unlimited number of derived keys.

In a Key Management System (KMS), users of the KMS may interact with the KMS to obtain encryption keys. An example of operation between a KMS requester (e.g., a computing device, a user such as associated with a computing device, etc.) and a KMS unit (e.g., another computing device, a KMS service, etc.) is as follows:

1. The requester seeking to access a key sends a Key Access Request (KAR) to a KMS unit, the request can include any one or more of:

- a. a requester identifier;
- b. a root key identifier;
- c. a sub-key identifier;
- d. authenticating information (e.g., credentials such as a password, a token, a response to a challenge, a signature, a digital certificate, etc.); and/or
- e. a challenge to the KMS unit (e.g., for the KMS unit to prove its identity or correctness of operation to the requester).

2. The KMS unit performs validation of the request, including any one or more of:

- a. Ensuring the credentials are correct for the requester identifier; and/or
- b. Ensuring the requester is authorized to access a key derived from the given root key identifier.

3. If not authorized, the KMS unit returns an error response and may create an audit log of the failure or take

other corrective actions. If the request is authorized, the KMS unit proceeds to the next step.

4. The KMS unit processes the access request, by using the appropriate root key (either the one indicated in the request, or by inferring it from other information, such as the requester identifier) together with the sub-key identifier to compute a sub-key. For example, when using a PRF to derive a sub-key, the KMS unit may compute that subkey S, as $S = \text{PRF}(\text{root-key}, \text{sub-key identifier})$. The KMS unit may create an audit log of the successful access request. It then proceeds to the next step.

5. If a challenge was provided by the requester to the KMS unit, the KMS unit generates a response to the challenge (e.g., a question, and a response to that question that compares favorably with the question)

6. The KMS unit returns a response to the requester including the sub-key and a challenge if one was generated

7. The requester validates the response to the challenge (if provided), and if it is valid, proceeds to use the sub-key (e.g., to perform encryption or decryption operations).

One downside to using a PRF in this way is that the KMS unit learns all the sub-keys returned to requesters, as the KMS unit computes the PRF and sees the input and output of the function. This makes the KMS service a central point of compromise for all the keys used by all the requesters in the system.

Such novel solutions as presented herein provides for applying a function known as an Oblivious Pseudorandom Function (OPRF). This can be enable secure access of keys by requesters from the KMS without the KMS being able to observe the keys and/or sub-keys that are requested and returned.

An OPRF works as follows. It takes two inputs:

1. an OPRF key “K” (e.g., an OPRF secret)
2. an arbitrary message “M” (e.g., a key ID, a label, a user-provided identifier, etc.)

From these two inputs, the OPRF also returns a pseudorandom output. However, unlike the PRF, the OPRF is computed by two parties (e.g., the requester and the KMS). The first party supplies the input “M”, while the second party supplies the OPRF key “K”. Only the first party receives (or can learn) the output of the OPRF. In the course of the computation, the first party does not learn any information about “K”. There are multiple implementations of OPRFs, including ones based on asymmetric encryption algorithms, RSA, blind signatures, Diffie-Hellman exponentiation, Elliptic Curve scalar multiplication, homomorphic encryption, and others. The general principle upon which OPRFs operate is that the first party obscures or “blinds” the input message “M” into a form which is meaningless to the second party before sending that second party. The second party then operates upon the input with a certain function that takes both the blinded input “B” along with the OPRF key “K” to produce a “blinded output” which is not the final output, but which is sent from the second party to the first party. The first party, with knowledge of how the original input was obscured, can then determine how to undo the impact of that operation from the blinded output, and recover the OPRF output. Because the reversal of the blinding is done by the first party, the second party never learns the OPRF output.

Taking the properties of the OPRF, and the design of the KMS described above together, the two may be merged to form a KMS which has superior security properties when compared to that which is provided in the prior art. This is done by substituting the PRF with an OPRF, and by having the requester perform some additional pre-processing of the

request and some post-processing of the response. The workflow with for an interaction with a KMS based on an OPRF might be as follows:

1. The requester obscures one of the inputs to a key derivation function, for example, a sub-key identifier, by using an appropriate blinding function for the OPRF that is used by the KMS unit. This produces a blinded-input "B". In some examples, the size of the blinding key is same as the size of the input provided thereto. For example, if the input is X bits or bytes in length, then the blinding key X bits or bytes in length (where X is a positive integer).

2. The requester seeking to access a key sends an Oblivious Key Access Request (OKAR) to a KMS unit, the request can include any one or more of:

- a. a requester identifier;
- b. a root key identifier (e.g., additional information to reference a specific OPRF key, e.g., a specific OPRF secret);
- c. a blinded input B (e.g., $B = \text{BlindingFunction}(\text{sub-key identifier})$);
- d. authenticating information (e.g., credentials such as a password, a token, a response to a challenge, a signature, a digital certificate, etc.); and/or
- e. A challenge to the KMS unit (for the KMS unit to prove its identity or correctness of operation)

3. The KMS unit performs validation of the request, including any one or more of:

- a. Ensuring the credentials are correct for the requester identifier; and/or
- b. Ensuring the requester is authorized to access a key derived from the given root key identifier.

4. If not authorized, the KMS unit may be configured to return an error response and may create an audit log of the failure or take other corrective actions. If the request is authorized, the KMS unit proceeds to the next step.

5. The KMS unit processes the access request, by using the appropriate root key (either the one indicated in the request, or by inferring it from other information, such as the requester identifier) together with the blinded input to compute a blinded sub-key. For example, when using an OPRF to derive a blinded sub-key, the KMS unit may compute that blinded subkey S, as $S = \text{OPRF}(\text{root-key}, B)$. The KMS unit may create an audit log of the successful access request. It then proceeds to the next step.

6. If a challenge was provided by the requester to the KMS unit, the KMS unit generates a response to the challenge.

7. The KMS unit returns a response to the requester including the blinded sub-key and a challenge if one was generated.

8. The requester validates the response to the challenge (if provided), and if it is valid, proceeds to unblind the sub-key using the appropriate function to unobscure the blinded sub-key and recover the OPRF output.

9. The requester uses the OPRF output as the key or to derive a key and then may perform encryption or decryption operations with that key.

In this manner, the KMS unit no longer sees the keys, and if the KMS unit cannot determine, predict, or guess the original unblinded sub-key identifiers, it has no capacity to determine any of the keys the requester receives.

FIG. 4A is a diagram illustrating an embodiment 401 of one or more communication systems supporting key protect with obliviousness according to various embodiments of the present invention. A service 410 (e.g., a cloud storage service, such as one that stores encrypted data 410a and Data Encryption Key (DEK) identifiers (IDs) 410b) and a key protect instance 410 (e.g., a Hardware Security Module

(HSM) 420b, such as one that stores one or more keys 420a). The service 410 blinds a DEK ID to generate $B(\text{DEK_ID})$. Then, the service 410 transmits get DEK request ($\text{getDEK}(\text{CDK_ID}, B(\text{DEK_ID}))$) to the key protect instance 410. The key protect instance 410 processes the get DEK request ($\text{getDEK}(\text{CDK_ID}, B(\text{DEK_ID}))$) and returned a blinded key $B(\text{DEK})$.

In this implementation, there is no information about a Data Encryption Key (DEK) that is exposed by the exchange between the service 410 (e.g., a cloud storage service) and the key protect instance 410 (e.g., a Hardware Security Module (HSM)). This oblivious implemented architecture's security, unlike prior art approaches in which various components and signals are vulnerable to interception during transmission, remains secure against adversaries with unbounded computing power as no useful information is revealed through the exchange. There is no information that is vulnerable to be intercepted during this process.

With the post-quantum security of obliviousness, the exchanged messages reveal no information to the attacker. Note that even when implemented, a Transport Layer Security's (TLS's) confidentiality is made superfluous given the security provided by the novel implementation as described herein. A breached oblivious key protect instance 420 would not endanger data keys, assuming the key ids are unknown to the attacker. Data Keys (Data Encryption Keys (DEKs)) only exist and available within the boundary of the service or user recovering the key. The only change on the service side is that instead of storing wrapped DEKs, the system would simply store the DEK IDs.

FIG. 4B is a diagram illustrating an embodiment 402 of one or more communication systems supporting Hardware Security Module (HSM) integration according to various embodiments of the present invention. An already-implemented HSM that supports the math used in accordance with OPRFs may be readily configured to support the novel implementation as presented herein. For example, such math that is based on an Elliptic Curve operation (e.g., EC Diffie-Hellman) may be used to support such OPRFs as described herein.

For example, when the math of OPRFs is fully supported by an existing HSMs, then Customer Root Keys (CRKs) can remain within HSMs at all times and will never be exposed to the host's memory.

Referring to the diagram, an application API DeriveKey operate is supported by client 430 based on a Data Encryption Key (DEK) identifier (ID) (DEK_ID) that undergoes blinding to generate a blinded data encryption key identifier, $B(\text{DEK})$ that is processed via the service 410 (e.g., cloud storage service) based on service-side HTTP REST API via the key protect instance 420 (e.g., HSM) as follows:

$$\text{Blinded}(\text{DEK}) \text{ or } B(\text{DEK}) = \text{DeriveKey}(\text{CRK_ID}, \text{Blinded}(\text{DEK_ID}))$$

This method takes two inputs, the customer root key identifier, CRK_ID, and the blinded data encryption key identifier, $\text{Blinded}(\text{DEK_ID})$. It returns a blinded data encryption key, $\text{Blinded}(\text{DEK})$ or $B(\text{DEK})$.

With respect to the Client-side SDK API, the Data Encryption Key (DEK) is returned as follows:

$$\text{DEK} = \text{DeriveKey}(\text{CRK_ID}, \text{DEK_ID})$$

This method takes the CRK id and the DEK id. The software development kit (SDK) code handles all blinding and de-blinding internally. It returns the DEK.

In some examples, to get obliviousness as a property, it may require some client-side preprocessing before invoking

21

the server's API, such as followed by some post-processing of the server's result. For example, this may be done to perform the blinding and de-blinding. A client-SDK would hide all of this from the end user and present a basic interface that takes the CRK and DEK IDs and returns the corresponding DEK.

In addition, note that multi-tenancy may also be supported such that different tenants supply different CRK_IDs that corresponds to a different OPRF key (e.g., different OPRF secret).

Note also that alternative, optional, and/or additional REST (RESTful, (representational state transfer)) API may be used as follows:

```
(CRVK-Certificate)=GetRootVerificationKey
(CRK_ID)
```

This method takes the CRK_ID and returns a certificate for the "Customer Root Verification Key" (CRVK) corresponding to the CRK.

The certificate binds the CRKV to the CRK_ID with a digital signature

The CRVK can be used to prove returned keys are correctly computed

```
Blinded(DEK),Response=DeriveAndVerifyKey
(CRK_ID,Blinded(DEK_ID),Challenge)
```

This method takes three inputs, the CRK_ID, the blinded DEK_ID, and a specially crafted "challenge". It enables the client to certify that the blinded key was computed correctly and using the correct CRK. This protects against MITM ("man in the middle," such as a middling device, etc.) attacks, server errors, defects, and memory corruption. Normally such a corruption would result in data loss such as based on encryption with a bad key. Note also that such novel solutions as presented herein can allow a client to validate that the KMS provided the correct key.

FIG. 4C is a diagram illustrating an embodiment 403 of key hierarchies as may be used in accordance with a KMS according to various embodiments of the present invention. This diagram includes a hierarchy that includes a root key at a top parent level, then development, sales, and operations in a lower child level, and then roadmap and sales at a child level below development, customers and leads at a child level below sales, and payroll at a child level below operations. Each respective lower level in the hierarchy is based on any encrypted by the key associated with one or more upper levels. This enables hierarchical business cases, e.g., "Root Key" encrypts "Development Key" encrypts "Roadmap Key". For another example, "Root Key" encrypts "Sales Key" encrypts "Leads Key". Access to a parent level key grants access to lower level child keys. Note that only knowing a key not directly in the lineage doesn't allow for access to other keys not in that lineage. For example, having "Development Key" wouldn't grant access to "Payroll".

Note that implementing a hierarchy of keys requires multiple levels of wrapping. If the hierarchy is deep, this can potentially introduce performance and scaling concerns. For each level of depth in the hierarchy, the KMS may need to import another key from the database and perform another HSM operation. Note that such novel solutions as described herein with respect to servicing and operating a KMS may be applied and applicable to any type of key hierarchy including system that include only one level therein or N levels therein (where N is a positive integer greater than or equal to 2).

In one embodiment, a variation of OPRF called Threshold Oblivious Pseudorandom Function (T-OPRF) adds thresh-

22

old security to OPRF. A Threshold Oblivious Pseudorandom Function (T-OPRF) is an OPRF computed not by a single server, but by some quorum of servers, none of which has knowledge of the Key. In practical terms, this enables an OPRF system having no single point of failure and no single point of compromise. Computing the T-OPRF requires two steps, first the evaluation of the T-OPRF by each server in the quorum to produce a partial output, and second the combination of these partial outputs to yield the final output, which again is learned only by the client, regardless of who performs this combination. In any T-OPRF, there is the total number of servers "N" and the minimum quorum size "T", which defines a "T-of-N" threshold scheme:

```
PartialOutputi=OPRF(PartialKeyi,Input)
```

```
Output=Combine(PartialOutput1,PartialOutput2,...
,PartialOutputT)
```

FIG. 5 is a schematic block diagram of an Threshold Oblivious Pseudorandom Function (T-OPRF) used in the process of producing encryption keys. OPRF systems can, in some instances, have a single-point-of-compromise (SPOC) if there are one or more locations which hold the OPRF key and use it to compute blinded OPRF outputs from the blinded inputs. The SPOC represents a potential security risk as an attacker that breaches this single location may, in some instances, be able to obtain the OPRF secret.

The construction of many OPRF algorithms, however, enables "thresholdizing" the OPRF key. This means a single OPRF key may be split into "N" distinct "fragments" (e.g., shown as partial keys 501-1 thru 501-N) and each fragment can be maintained at a corresponding distinct location (e.g., in a different Key Management System (KMS) unit (shown as 220-1 thru 220-N)). Further, the thresholdizing means that any threshold number "T" (where 1<=T<=N) or more results are obtained from T or more KMS units each holding a distinct fragment (partial key) of the OPRF key, and each participating in an OPRF operation, can be combined to yield a blinded output of the OPRF (discussed in greater detail in association with FIG. 6).

In addition, no subset of fragments of the OPRF key less than the threshold T can be used to compute the OPRF. When T is greater than one, there is no single point of compromise, and increasing T increases the number of locations that must be compromised to lead to an exposure of the OPRF key or otherwise enable an adversary to compute OPRF outputs.

Multiple constructions for threshold OPRFs exist, and include Blinded Threshold RSA Signatures, Diffie-Hellman exponentiation with "Shamir in the Exponent", Elliptic Curve point multiplication with Shamir in the Exponent. These constructions permit the OPRF secret to be divided into N fragments, with any T necessary and sufficient to assist in the determination of the OPRF output. As with conventional OPRFs, the participants holding the fragments (partial keys) and assisting in the computation learn no information about either the input or the output of the OPRF.

A requester seeking to access a key sends an Oblivious Key Access Request (OKAR) to at least a threshold number T of distinct KMS units. The requester may do this directly by itself sending the request to a threshold T number of distinct KMS units, or by sending a single request to an optional "intermediate layer" 502, with the intermediate layer then selecting at least a threshold T number of KMS units to send the request to. While shown as positioned between the cloud storage service 210 and the multiple KMS units for illustration purposes, the intermediate layer can be

functionally configured anywhere between the requester and the KMS units. By using an intermediate layer, existing client's protocols and APIs need not be changed when upgrading from a KMS based on an OPRF to one based on a Threshold OPRF. The protocol can appear the same in terms of the request and response seen by the requester when using an intermediate layer with a T-OPRF KMS, as with a requester communicating directly to an OPRF KMS.

FIG. 6 is another embodiment of a block diagram of a Threshold Oblivious Pseudorandom Function (T-OPRF) used in the process of producing encryption keys according to various embodiments of the present invention. A Key Management System based on a Threshold OPRF (T-OPRF) operates as follows:

1. N distinct KMS services/units are deployed (220-1 through 220-N)

2. N distinct fragments (partial keys PK-1 through PK-N (606)) of the OPRF key (608) are installed to each distinct KMS unit, using a distribution protocol

- a. In one distribution protocol embodiment, the OPRF key is generated at one location and then fragments of it are computed and sent separately to one of the KMS units
- b. In an alternate distribution protocol embodiment, the KMS units cooperate to derive consistent fragments of an OPRF key without the OPRF key ever existing in any location and without any KMS unit being able to determine it.

After the installation of the OPRF key fragments to the set of KMS units, a requester may obtain keys from the T-OPRF KMS as follows:

1. The requester obscures one of the inputs 602 to a key derivation function, for example, a sub-key identifier, by using an appropriate blinding function (blinding key 604) for the OPRF that is used by the KMS unit. This produces a blinded-input "B"

2. The requester seeking to access a key sends an Oblivious Key Access Request (OKAR) to at least a threshold number of distinct KMS units. The requester may do this either by itself sending the request to a threshold T number of distinct KMS units, or by sending a single request to an "intermediate layer", and the intermediate layer then selects at least a threshold T number of KMS units to send the request to. The request generated by the requester can include:

- a. a requester identifier
- b. a root key identifier—(Additional information to reference a specific OPRF key)
- c. a blinded input B (e.g. $B = \text{BlindingFunction}(\text{sub-key identifier})$)
- d. authenticating information (e.g. credentials such as a password, a token, a response to a challenge, a signature, a digital certificate, etc.)
- e. A challenge to the KMS unit (for the KMS unit to prove its identity or correctness of operation)

3. Each KMS unit that receives a request performs validation of the request, including:

- a. ensuring the credentials are correct for the requester identifier
- b. ensuring the requester is authorized to access a key derived from the given root key identifier

4. If not authorized, the KMS unit returns an error response and may create an audit log of the failure or take other corrective actions. If the request is authorized, the KMS unit proceeds to the next step.

5. The KMS unit processes the access request, by using the appropriate root key (either the one indicated in the request, or by inferring it from other information, such as the

requester identifier) together with the blinded input to compute a blinded sub-key (blinded partial key 610). For example, when using an OPRF 608 to derive a blinded sub-key 610, the KMS unit may compute that blinded subkey S, as $S = \text{OPRF}(\text{root-key-fragment}, B)$. The KMS unit may create an audit log of the successful access request. It then proceeds to the next step.

6. If a challenge was provided by the requester to the KMS unit, the KMS unit generates a response to the challenge.

7. The KMS unit returns a response to the requester including the blinded sub-key and a challenge if one was generated, or if the requester used an intermediate layer, the KMS unit returns the response to the intermediate layer.

8. The receiver of the response (either the requester or the intermediate layer) validates the response to the challenge (if provided), and if it is valid, and there are at least a threshold number of returned validated responses proceeds to combine 612 the threshold or more responses through the appropriate thresholding algorithm to obtain the "blinded OPRF output" 614. If there is less than a threshold number of validated responses, the receiver may request additional responses from other KMS units not involved in the initial operation. If the receiver is unable to acquire a threshold number of validated responses it will return an error to the requester. If the receiver was an intermediate layer, it returns the "blinded OPRF output" to the requester.

9. The requester unblinds the blinded OPRF output using the appropriate function (e.g., blinding key 604) to unobscure the blinded OPRF output and recover the OPRF output.

10. The requester uses the OPRF output as the sub-key or to derive a key and then may perform encryption or decryption operations with that key.

In this manner, the KMS unit not only never sees the keys (by virtue of the OPRF), but no single set of KMS units less than the threshold, even attempt to determine, the resulting sub-keys even if the original unblinded sub-key identifiers are known. This makes the security of the KMS based on a T-OPRF greater than what can be achieved by a KMS based on an OPRF (without thresholding).

FIG. 7 is a diagram illustrating an embodiment of a method for execution by one or more computing devices according to various embodiments of the present invention.

The method 700 begins, in step 702, by dividing an Oblivious Pseudorandom Function (OPRF) key into key fragments to generate a plurality of N partial keys. The method continues, in step 704, by distributing, via a communication system, a respective one of the plurality of N partial keys to a corresponding plurality of N Key Management System (KMS) units, the plurality of N partial keys distributed according to a distribution protocol for storage in the N KMS units. The method continues, in step 706, by receiving, via the communication system and from a threshold number T of the KMS units, a plurality T of partial blinded keys, wherein the plurality T of partial blinded keys are based on processing of a value of a blinded key received by a respective KMS unit and a corresponding stored partial key of the N partial keys. The T partial blinded keys are created using processes previously described methods of FIG. 6 as well as those described in FIGS. 1-5.

The method continues, in step 708, by combining the plurality T of partial blinded keys into the blinded key. While illustrated as a high-level combining of partial blinded keys, one or more sub elements directed to previously described requester OKAR processing, validation, challenges, returning a threshold number T of responses

from the T KMS units, are considered part of KMS unit activity. For example, upon the access of a structured partial key from the OPRF-based KMS, the Key Management System will return the partial blinded output of the OPRF (which is pseudorandom) as well as the structured partial key parameters, a generating procedure, and optionally encrypted parameters. The requester of the key uses the values returned by the KMS to produce the structured key as follows:

1. The requester obtains the output of the OPRF (e.g., by unblinding the blinded OPRF output returned from the KMS)

2. The requester determines the complete set of parameters from:

a. The structured key parameters; and/or
b. When encrypted parameters are returned, by further deriving a key from the OPRF output to decrypt the encrypted parameters to obtain the decrypted parameters

3. Using the complete set of parameters, and the OPRF output, following the returned “Generating Procedure” to produce the Structured Key

The generating procedure may be something very straight forward, such as expanding the OPRF output to a specific length (e.g., using a mask generating function) and then taking the result modulo a prime modulus specific to the Diffie-Hellman or Elliptic Curve parameters, or it might be something more complex, such as using the OPRF output to seed a pseudorandom number generator (PRNG), using the PRNG to generate two large prime numbers, and using these two large prime numbers to produce an RSA modulus and then from the RSA modulus and the large prime numbers produce an RSA key pair. Note however, that a process like this can be resource intensive, as generating RSA keys and other structured keys can be expensive and time consuming. For these cases, it may be more efficient to utilize encrypted parameters. For example, the encrypted parameters might be the two large primes from which an RSA key can be derived, or it might even be the structured key itself. In this case, the generating procedure might specify how to derive a symmetric encryption key from the OPRF output, and then use this key to decrypt the encrypted parameters, and from these encrypted parameters, produce the structured key.

The method continues in step **710** by processing the blinded key based on the blinding key in accordance with an OPRF unblinding operation to generate a key and, in step **712**, accessing secure information based on the key. For example, the key can be used to encrypt, or decrypt data stored in the cloud-based system.

FIG. **8** is a diagram illustrating an embodiment of a method for execution by one or more computing devices according to various embodiments of the present invention. For N distinct KMS services/units deployed (**220-1** through **220-N**), method **800** begins, in step **802**, by transmitting, for installation, N distinct partial keys of an Oblivious Pseudorandom Function (OPRF) key to each of N distinct KMS units, using a distribution protocol. In one distribution protocol embodiment, the OPRF key is generated at one location and then fragments of it are computed and sent separately to one of the KMS units. In an alternate distribution protocol embodiment, the KMS units cooperate to derive consistent fragments of an OPRF key without the OPRF key ever existing in any location and without any KMS unit being able to determine it.

The method continues, in step **804**, by transmitting an Oblivious Key Access Request (OKAR) to at least a threshold number T of the N distinct KMS units. A requester may do this either by itself sending the request to a threshold T

number of distinct KMS units, or by sending a single request to an “intermediate layer”, and the intermediate layer then selects at least a threshold T number of KMS units to send the request to. The request generated by the requester can include:

- a. a requester identifier
- b. a root key identifier—(Additional information to reference a specific OPRF key)
- c. a blinded input B (e.g. $B = \text{BlindingFunction}(\text{sub-key identifier})$)
- d. authenticating information (e.g. credentials such as a password, a token, a response to a challenge, a signature, a digital certificate, etc.)
- e. A challenge to the KMS unit (for the KMS unit to prove its identity or correctness of operation)

The method continues, in step **806**, by performing a validation of the OKAR for at least a threshold number T of the N distinct KMS units. Each KMS unit that receives a request performs validation of the request, including: ensuring the credentials are correct for the requester identifier and/or ensuring the requester is authorized to access a key derived from the given root key identifier. If not authorized, the KMS unit returns an error response and may create an audit log of the failure or take other corrective actions. If the request is authorized, the KMS unit proceeds to the next step.

The method continues, in step **808**, for a validated OKAR, by receiving an OKAR response from a corresponding one of the T number of the KMS units, the OKAR response uses a root partial key combined with a blinded input to compute a blinded partial-key. For example, the KMS unit processes the access request, by using the appropriate root key (either the one indicated in the request, or by inferring it from other information, such as the requester identifier) together with the blinded input to compute a blinded sub-key (blinded partial key). For example, when using an OPRF to derive a blinded sub-key, the KMS unit may compute that blinded subkey S, as $S = \text{OPRF}(\text{root-key-fragment}, B)$. The KMS unit may create an audit log of successful/unsuccessful access requests.

The method continues, in step **810**, when at least a threshold number T of validated OKAR responses are received, by combining the at least a threshold number T of validated OKAR responses to obtain a blinded OPRF output. If a challenge was provided by the requester to the KMS unit, the KMS unit generates a response to the challenge and returns a response to the requester including the blinded sub-key and a challenge if one was generated, or if the requester used an intermediate layer, the KMS unit returns the response to the intermediate layer. The receiver of the response (either the requester or the intermediate layer) validates the response to the challenge (if provided), and if it is valid, and there are at least a threshold number of returned validated responses proceeds to combine the threshold or more responses through the appropriate thresholding algorithm to obtain the “blinded OPRF output”. If there is less than a threshold number of validated responses, the receiver may request additional responses from other KMS units not involved in the initial operation. If the receiver is unable to acquire a threshold number of validated responses it will return an error to the requester. If the receiver was an intermediate layer, it returns the “blinded OPRF output” to the requester.

The method continues in step **812** by unblinding the blinded OPRF output to recover an OPRF output to be used as a key. The requester unblinds the blinded OPRF output using the appropriate function (e.g., blinding key **604**) to

unobscure the blinded OPRF output and recover the OPRF output. The requester uses the OPRF output as the sub-key or to derive a key and, in step 814, performs encryption or decryption operations with the key.

In this manner, the KMS unit not only never sees the keys (by virtue of the OPRF), but no single set of KMS units less than the threshold, even attempt to determine, the resulting sub-keys even if the original unblinded sub-key identifiers are known. This makes the security of the KMS based on a T-OPRF greater than what can be achieved by a KMS based on an OPRF (without thresholding).

In some examples, the input value is unknown to the other computing device. Also, in certain specific examples, the input value includes a key identifier that is associated with the key. Also, in some examples, the key is unknown to the other computing device. In addition, in certain specific examples, the key includes a Data Encryption Key (DEK) or a Key Encryption Key (KEK). Note that the OPRF secret is unknown to the computing device and is based on a Customer Root Key (CRK) that is associated with the computing device.

With respect to solutions using an OPRF such as described herein, note that the output of Pseudorandom Functions and Oblivious Pseudorandom Functions may be uncontrollable and appear completely random. In some examples, this can present a problem for key types that require a certain structures, formats, or properties (e.g., "Structured Keys"). For example, an RSA (Rivest-Shamir-Adleman) key consists of a set of numbers bearing a specific relation to each other. Similarly, an Elliptic Curve key consists of a number on a specific range as defined by the curve for which it is associated. The output of an OPRF is not guaranteed to produce outputs that will match the desired properties or required structures for keys. To support keys with defined structures, an OPRF-based Key Management System can be extended to support structured keys as described below.

For each structured partial key managed by the OPRF-based Key Management System, the Key Management System (KMS) will maintain:

1. Structured partial Key Parameters (one or more of: the key type and algorithm (RSA, DH, ElGamal, Elliptic Curve, etc.), key size (256-bits, 512-bits, 1024-bits, 2048-bits, etc.), modulus, generator, curve parameters)

2. A Generating Procedure (a deterministic algorithm to produce the structured partial key from the output of the OPRF, the structured partial key parameters, and optionally the encrypted parameters (if present)).

3. Optionally: Encrypted Parameters (may contain secret information or parameters encrypted under a partial key which can be produced from the OPRF output).

FIG. 9 depicts a cloud computing environment 900 according to various embodiments of the present invention. FIG. 9 presents an illustrative cloud computing environment 50. As shown, cloud computing environment 50 includes one or more cloud computing nodes 10 with which local computing devices used by cloud consumers, such as, for example, personal digital assistant (PDA) or cellular telephone 54A, desktop computer 54B, laptop computer 54C, and/or automobile computer system 54N may communicate. Nodes 10 may communicate with one another. They may be grouped (not shown) physically or virtually, in one or more networks, such as Private, Community, Public, or Hybrid clouds as described hereinabove, or a combination thereof. This allows cloud computing environment 50 to offer infrastructure, platforms and/or software as services for which a cloud consumer does not need to maintain resources on a

local computing device. It is understood that the types of computing devices 54A-N shown in FIG. 9 are intended to be illustrative only and that computing nodes 10 and cloud computing environment 50 can communicate with any type of computerized device over any type of network and/or network addressable connection (e.g., using a web browser).

It is to be understood that although this disclosure includes a detailed description on cloud computing, implementation of the teachings recited herein are not limited to a cloud computing environment. Rather, embodiments of the present invention are capable of being implemented in conjunction with any other type of computing environment now known or later developed.

Cloud computing is a model of service delivery for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, network bandwidth, servers, processing, memory, storage, applications, virtual machines, and services) that can be rapidly provisioned and released with minimal management effort or interaction with a provider of the service. This cloud model may include at least five characteristics, at least three service models, and at least four deployment models.

Characteristics are as follows:

On-demand self-service: a cloud consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with the service's provider.

Broad network access: capabilities are available over a network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

Resource pooling: the provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to demand. There is a sense of location independence in that the consumer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter).

Rapid elasticity: capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured service: cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

Service Models are as follows:

Software as a Service (SaaS): the capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based e-mail). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Platform as a Service (PaaS): the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider.

The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Infrastructure as a Service (IaaS): the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

Deployment Models are as follows:

Private cloud: the cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.

Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

Public cloud: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Hybrid cloud: the cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

A cloud computing environment is service oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability. At the heart of cloud computing is an infrastructure that includes a network of interconnected nodes.

FIG. 10 depicts abstraction model layers 1000 according to various embodiments of the present invention. Referring now to FIG. 10, a set of functional abstraction layers provided by cloud computing environment 50 (FIG. 9) is shown. It should be understood in advance that the components, layers, and functions shown in FIG. 10 are intended to be illustrative only and embodiments of the invention are not limited thereto. As depicted, the following layers and corresponding functions are provided:

Hardware and software layer 60 includes hardware and software components. Examples of hardware components include: mainframes 61; RISC (Reduced Instruction Set Computer) architecture-based servers 62; servers 63; blade servers 64; storage devices 65; and networks and networking components 66. In some embodiments, software components include network application server software 67 and database software 68. In some embodiments, one or more hardware components can be implemented by utilizing a computing device as is known in the art.

Virtualization layer 70 provides an abstraction layer from which the following examples of virtual entities may be provided: virtual servers 71; virtual storage 72; virtual networks 73, including virtual private networks; virtual applications and operating systems 74; and virtual clients 75.

In one example, management layer 80 may provide the functions described below. Resource provisioning 81 provides dynamic procurement of computing resources and other resources that are utilized to perform tasks within the

cloud computing environment. Metering and Pricing 82 provide cost tracking as resources are utilized within the cloud computing environment, and billing or invoicing for consumption of these resources. In one example, these resources may include application software licenses. Security provides identity verification for cloud consumers and tasks, as well as protection for data and other resources. User portal 83 provides access to the cloud computing environment for consumers and system administrators. Service level management 84 provides cloud computing resource allocation and management such that required service levels are met. Service Level Agreement (SLA) planning and fulfillment 85 provide pre-arrangement for, and procurement of, cloud computing resources for which a future requirement is anticipated in accordance with an SLA.

Workloads layer 90 provides examples of functionality for which the cloud computing environment may be utilized. Examples of workloads and functions which may be provided from this layer include: mapping and navigation 91; software development and lifecycle management 92; virtual classroom education delivery 93; data analytics processing 94; transaction processing 95; and security, encryption, and key management related processing, and operations 96.

As may be used herein, the terms “substantially” and “approximately” provides an industry-accepted tolerance for its corresponding term and/or relativity between items. Such an industry-accepted tolerance ranges from less than one percent to fifty percent and corresponds to, but is not limited to, component values, integrated circuit process variations, temperature variations, rise and fall times, and/or thermal noise. Such relativity between items ranges from a difference of a few percent to magnitude differences. As may also be used herein, the term(s) “configured to”, “operably coupled to”, “coupled to”, and/or “coupling” includes direct coupling between items and/or indirect coupling between items via an intervening item (e.g., an item includes, but is not limited to, a component, an element, a circuit, and/or a module) where, for an example of indirect coupling, the intervening item does not modify the information of a signal but may adjust its current level, voltage level, and/or power level. As may further be used herein, inferred coupling (i.e., where one element is coupled to another element by inference) includes direct and indirect coupling between two items in the same manner as “coupled to”. As may even further be used herein, the term “configured to”, “operable to”, “coupled to”, or “operably coupled to” indicates that an item includes one or more of power connections, input(s), output(s), etc., to perform, when activated, one or more its corresponding functions and may further include inferred coupling to one or more other items. As may still further be used herein, the term “associated with”, includes direct and/or indirect coupling of separate items and/or one item being embedded within another item.

As may be used herein, the term “compares favorably”, indicates that a comparison between two or more items, signals, etc., provides a desired relationship. For example, when the desired relationship is that signal 1 has a greater magnitude than signal 2, a favorable comparison may be achieved when the magnitude of signal 1 is greater than that of signal 2 or when the magnitude of signal 2 is less than that of signal 1. As may be used herein, the term “compares unfavorably”, indicates that a comparison between two or more items, signals, etc., fails to provide the desired relationship.

As may also be used herein, the terms “processing module”, “processing circuit”, “processor”, and/or “processing unit” may be a single processing device or a plurality of

processing devices. Such a processing device may be a microprocessor, micro-controller, digital signal processor, microcomputer, central processing unit, field programmable gate array, programmable logic device, state machine, logic circuitry, analog circuitry, digital circuitry, and/or any device that manipulates signals (analog and/or digital) based on hard coding of the circuitry and/or operational instructions. The processing module, module, processing circuit, and/or processing unit may be, or further include, memory and/or an integrated memory element, which may be a single memory device, a plurality of memory devices, and/or embedded circuitry of another processing module, module, processing circuit, and/or processing unit. Such a memory device may be a read-only memory, random access memory, volatile memory, non-volatile memory, static memory, dynamic memory, flash memory, cache memory, and/or any device that stores digital information. Note that if the processing module, module, processing circuit, and/or processing unit includes more than one processing device, the processing devices may be centrally located (e.g., directly coupled together via a wired and/or wireless bus structure) or may be distributedly located (e.g., cloud computing via indirect coupling via a local area network and/or a wide area network). Further note that if the processing module, module, processing circuit, and/or processing unit implements one or more of its functions via a state machine, analog circuitry, digital circuitry, and/or logic circuitry, the memory and/or memory element storing the corresponding operational instructions may be embedded within, or external to, the circuitry comprising the state machine, analog circuitry, digital circuitry, and/or logic circuitry. Still further note that, the memory element may store, and the processing module, module, processing circuit, and/or processing unit executes, hard coded and/or operational instructions corresponding to at least some of the steps and/or functions illustrated in one or more of the Figures. Such a memory device or memory element can be included in an article of manufacture.

One or more embodiments have been described above with the aid of method steps illustrating the performance of specified functions and relationships thereof. The boundaries and sequence of these functional building blocks and method steps have been arbitrarily defined herein for convenience of description. Alternate boundaries and sequences can be defined so long as the specified functions and relationships are appropriately performed. Any such alternate boundaries or sequences are thus within the scope and spirit of the claims. Further, the boundaries of these functional building blocks have been arbitrarily defined for convenience of description. Alternate boundaries could be defined as long as the certain significant functions are appropriately performed. Similarly, flow diagram blocks may also have been arbitrarily defined herein to illustrate certain significant functionality.

To the extent used, the flow diagram block boundaries and sequence could have been defined otherwise and still perform the certain significant functionality. Such alternate definitions of both functional building blocks and flow diagram blocks and sequences are thus within the scope and spirit of the claims. One of average skill in the art will also recognize that the functional building blocks, and other illustrative blocks, modules and components herein, can be implemented as illustrated or by discrete components, application specific integrated circuits, processors executing appropriate software and the like or any combination thereof.

In addition, a flow diagram may include a “start” and/or “continue” indication. The “start” and “continue” indica-

tions reflect that the steps presented can optionally be incorporated in or otherwise used in conjunction with other routines. In this context, “start” indicates the beginning of the first step presented and may be preceded by other activities not specifically shown. Further, the “continue” indication reflects that the steps presented may be performed multiple times and/or may be succeeded by other activities not specifically shown. Further, while a flow diagram indicates a particular ordering of steps, other orderings are likewise possible provided that the principles of causality are maintained.

The one or more embodiments are used herein to illustrate one or more aspects, one or more features, one or more concepts, and/or one or more examples. A physical embodiment of an apparatus, an article of manufacture, a machine, and/or of a process may include one or more of the aspects, features, concepts, examples, etc. described with reference to one or more of the embodiments discussed herein. Further, from figure to figure, the embodiments may incorporate the same or similarly named functions, steps, modules, etc. that may use the same or different reference numbers and, as such, the functions, steps, modules, etc. may be the same or similar functions, steps, modules, etc. or different ones.

Unless specifically stated to the contra, signals to, from, and/or between elements in a figure of any of the figures presented herein may be analog or digital, continuous time or discrete time, and single-ended or differential. For instance, if a signal path is shown as a single-ended path, it also represents a differential signal path. Similarly, if a signal path is shown as a differential path, it also represents a single-ended signal path. While one or more particular architectures are described herein, other architectures can likewise be implemented that use one or more data buses not expressly shown, direct connectivity between elements, and/or indirect coupling between other elements as recognized by one of average skill in the art.

The term “module” is used in the description of one or more of the embodiments. A module implements one or more functions via a device such as a processor or other processing device or other hardware that may include or operate in association with a memory that stores operational instructions. A module may operate independently and/or in conjunction with software and/or firmware. As also used herein, a module may contain one or more sub-modules, each of which may be one or more modules.

The present invention may be a system, a method, and/or a computer program product at any possible technical detail level of integration. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

As may further be used herein, a computer readable memory includes one or more memory elements. A memory element may be a separate memory device, multiple memory devices, or a set of memory locations within a memory device. Such a memory device may be a read-only memory, random access memory, volatile memory, non-volatile memory, static memory, dynamic memory, flash memory, cache memory, and/or any device that stores digital information. The memory device may be in a form a solid-state memory, a hard drive memory, cloud memory, thumb drive, server memory, computing device memory, and/or other physical medium for storing digital information.

The computer readable storage medium can be a tangible device that can retain and store instructions for use by an

instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, configuration data for integrated circuitry, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++, or the like, and procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of

methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

These computer readable program instructions may be provided to a processor of a general-purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the blocks may occur out of the order noted in the Figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

While particular combinations of various functions and features of the one or more embodiments have been expressly described herein, other combinations of these features and functions are likewise possible. The present disclosure is not limited by the particular examples disclosed herein and expressly incorporates these other combinations.

What is claimed is:

1. A computing device comprising:

an interface configured to interface and communicate with a communication system; memory that stores operational instructions; and

processing circuitry operably coupled to the interface and to the memory, wherein the processing circuitry is configured to execute the operational instructions to:

35

divide an Oblivious Pseudorandom Function (OPRF) key to generate a plurality of N partial keys;

distribute, via the communication system, a respective one of the plurality of N partial keys to a corresponding plurality of N Key Management System (KMS) units, the plurality of N partial keys distributed according to a distribution protocol for storage in the N KMS units;

receive, via the communication system and from a threshold number T of the KMS units, a plurality T of partial blinded keys, wherein the plurality T of partial blinded keys are based on processing of a value of a blinded key received by a respective KMS unit and a corresponding stored partial key of the N partial keys;

combine the plurality T of partial blinded keys into the blinded key;

process the blinded key based on the blinding key in accordance with an OPRF unblinding operation to generate a key; and

access secure information based on the key.

2. The computing device of claim 1, wherein the distribution protocol includes generating the OPRF key at one location with the N partial keys computed and sent separately to one of the plurality of N KMS units.

3. The computing device of claim 1, wherein the distribution protocol includes the plurality of N KMS units cooperating to derive the N partial keys of the OPRF key without the OPRF key existing in any single location.

4. The computing device of claim 1, wherein the blinded key received by a respective KMS unit is included in a specific request to a threshold T number of distinct KMS units.

5. The computing device of claim 1, wherein the blinded key received by a respective KMS unit is included as a single request to an intermediate layer, with the intermediate layer then selecting at least a threshold T number of KMS units to send a specific request to.

6. The computing device of claim 1, wherein the computing device includes any of: a wireless smart phone, a cellular phone, a laptop, a personal digital assistant, a tablet, a personal computers (PC), a work station, a video game device, or another computing device including a Hardware Security Module (HSM).

7. The computing device of claim 1, wherein the communication system includes at least one of a wireless communication system, a wire lined communication system, a nonpublic intranet system, a public internet system, a local area network (LAN), a wireless local area network (WLAN), a wide area network (WAN), a satellite communication system, a fiberoptic communication system, or a mobile communication system.

8. A method for secure processing, the method comprising:

dividing an Oblivious Pseudorandom Function (OPRF) key to generate a plurality of N partial keys;

distributing, via a communication system, a respective one of the plurality of N partial keys to a corresponding plurality of N Key Management System (KMS) units, the plurality of N partial keys distributed according to a distribution protocol for storage in the N KMS units;

receiving, via the communication system and from a threshold number T of the KMS units, a plurality T of partial blinded keys, wherein the plurality T of partial blinded keys are based on processing of a value of a blinded key received by a respective KMS unit and a corresponding stored partial key of the N partial keys;

36

combining the plurality T of partial blinded keys into the blinded key;

processing the blinded key based on the blinding key in accordance with an OPRF unblinding operation to generate a key; and

accessing secure information based on the key.

9. The method of claim 8, wherein the distribution protocol includes generating the OPRF key at one location with the N partial keys computed and sent separately to one of the plurality of N KMS units.

10. The method of claim 8, wherein the distribution protocol includes the plurality of N KMS units cooperating to derive the N partial keys of the OPRF key without the OPRF key existing in any single location.

11. The method of claim 8, wherein the blinded key received by a respective KMS is included in a specific request to a threshold T number of distinct KMS units.

12. A system comprises:

an interface configured to interface and communicate with a communication system; memory that stores operational instructions; and

processing circuitry operably coupled to the interface and to the memory, wherein the processing circuitry is configured to execute the operational instructions to:

transmit, for installation, N distinct partial keys of an Oblivious Pseudorandom Function (OPRF) key to each of N distinct KMS units, using a distribution protocol;

transmit an Oblivious Key Access Request (OKAR) to at least a threshold number T of the N distinct KMS units;

perform a validation of the OKAR for at least a threshold number T of the N distinct KMS units;

for a validated OKAR:

receive an OKAR response from a corresponding one of the T number of the KMS units, the OKAR response using a root partial key combined with a blinded input to compute a blinded partial-key;

when at least a threshold number T of validated OKAR responses are received, combine the at least a threshold number T of validated OKAR responses to obtain a blinded OPRF output; and unblind the blinded OPRF output to recover an OPRF output to be used as a key; and

perform encryption or decryption operations with the key.

13. The system of claim 12 further comprising operational instructions to derive the N distinct partial keys based on a partial-key identifier.

14. The system of claim 13, wherein the Oblivious Key Access Request (OKAR) includes one or more of:

a requester identifier;

a root key identifier;

a blinded input B where $B = \text{BlindingFunction}(\text{partial-key identifier})$;

authenticating information, or a challenge to the KMS unit.

15. The system of claim 14, wherein the authenticating information includes one or more of:

credentials;

a password;

a token;

a response to a challenge;

a signature; or

a digital certificate.

16. The system of claim 15, wherein the response to a challenge requires a KMS unit to prove its identity or correctness of operation.

17. The system of claim 15, wherein the validation of the OKAR includes any of: ensuring credentials are correct for a requester identifier or ensuring a requester is authorized to access a key derived from a given root key identifier.

18. The system of claim 12, wherein, when the validation 5 indicates that a KMS unit is not authorized, receiving an error response from the KMS unit.

19. The system of claim 12 further comprises, when using an OPRF to derive a blinded partial-key reflects the KMS unit computing a blinded partial-key S, as $S = \text{OPRF}(\text{rootkey-}$ 10 $\text{partial}, B)$.

20. The system of claim 12 further comprises, when a challenge is provided by a requester to the KMS unit, the KMS unit generating a response to the challenge and returning a response to the requester including the blinded partial- 15 key and the challenge.

* * * * *